# The Impossibility Theorems of Cartography

Gauss, Tissot, and the Topological Limits of the Sphere

Silvano Salvador

# Preface and Status of This Volume

This is a self-contained, stand-alone volume in the tradition of the cartographic distortion series (*No Map Is Best*), the Discrete Global Grid Systems monograph, and *The Retroazimuthal Problem*. Where those volumes examine specific projections, specific administrative datums, or a specific historical episode, this one examines the theorems that make all of them necessary: the proofs that certain properties a cartographer might want — simultaneous conformality and equal-area representation, a globally smooth field of fixed directions, an isometric flattening of the sphere — are not merely difficult to achieve, but impossible in a precise, provable sense.

**Epistemic Notice**

All four parts of this volume now contain at least a full first pass. Part I (Foundations) is complete and has been checked against the primary sources cited: Gauss's *Disquisitiones generales circa superficies curvas* (1827/1828) and Tissot's *Mémoire sur la représentation des surfaces et les projections des cartes géographiques* (1881). Part II (the four central impossibility theorems) is drafted in full; chapters 4, 6 and 7 are proved in full from material established earlier in the volume, while chapter 5 rests, for its optimality half specifically, on two independent but still secondary-sourced accounts of Milnor's 1969 proof rather than a primary-verified reading, as flagged in place. Part III (now seven historical case-study chapters, eight named constructions: Mercator/Lambert, chapter 14; Gauss/Krüger, chapter 16; discrete global grid systems, chapter 18; Robinson as the aphylactic/compromise case, chapter 19, confirmed aphylactic by direct computation rather than by design intent; Goode's homolosine as the dedicated interruption case, chapter 20, whose join parallel this volume re-derives and verifies numerically against the historical value, and whose two equal-area halves are now both proved in closed form rather than one proved and one cited; Fuller's Dymaxion projection as the dedicated aphylactic-*and*-interrupted case, chapter 21, closing the family-by-continuity grid's aphylactic-interrupted cell; Peirce's quincuncial projection, chapter 22, examined specifically for the remaining conformal-and-interrupted cell and found, by a from-scratch proof that composition of conformal maps is conformal, to be conformal but contiguous rather than interrupted — a sharpened, not closed, account of that last gap; and, closing that last gap in turn, this volume's own stereographic-icosahedral construction, chapter 23, assembled entirely from results already proved elsewhere in the volume rather than from any new cited material, conformal on every point of every one of its twenty gores and interrupted by the same combinatorial cutting chapters 20 and 21 used, closing the family-by-continuity grid's last cell by a constructed witness rather than a historical case) and Part IV (a bridge chapter formally connecting proposition 18.1 to chapter 6 via cone-point holonomy and the Poincaré–Hopf theorem, chapter 25, extended by a refining-sequence section, section 25.7, that establishes Hausdorff convergence of successive icosahedral Goldberg polyhedra to the round sphere, shows the discrete total-curvature identity needs no limiting argument for a genuinely flat-faced polyhedron (it is exact at every finite stage of refinement), and, in a computed subsection, section 25.7.1, reports the sixty pentagon-corner defects explicitly for one concrete construction (9.105861° each, against a 12° idealized estimate), proves the pentagon contribution to each is exactly 108° by symmetry, and discovers and measures — rather than merely allows for in principle — the non-planarity of the adjoining hexagonal faces responsible for the remaining convention-dependence; a further subsection, section 25.7.2, extends this computation to $n = 3, 4$ and to a second, face-flattening convention, finding the two conventions agree to within about a tenth of a percent (section 25.7.2) and, more unexpectedly, that the curvature this construction carries does not stay concentrated at the twelve original pentagons as refinement proceeds but spreads increasingly onto the hexagon-only vertices instead (section 25.7.2), leaving the precise asymptotic rate as a newly sharpened open item rather than a closed one, narrowed again by a subsequent least-squares fit extending the data to $n = 40$ (section 25.7.4), whose restricted large-$n$ slope, $-2.079 \pm 0.007$, sits closer to the conjectured $-2$ than earlier estimates without its own confidence interval quite containing it, narrowed once more by a further extension to $n = 50$ (section 25.7.5), whose tightest range yet ($n = 20$ to 50) gives slope $-2.049 \pm 0.005$, closer still but again short of $-2$; and by a companion analysis of the full defect distribution's relative shape rather than only its overall scale (section 25.7.6), whose three computed shape statistics plateau well short of the values exact uniformity would require, suggesting — without proving — that the discrete curvature settles toward a fixed non-constant profile on $S^2$ rather than flattening out entirely, a finding that reframes rather than answers the chapter's open question of what continuous object the discrete singularities converge to; a third subsection, section 25.7.48, tests a third, independent convention (centroid-fan triangulation) named but left untested in earlier drafts, finding it agrees pointwise with the first two even more closely than they agree with each other, while its whole-polyhedron total appears at first to depart measurably from the exact 720° the other two conventions match, a departure section 25.7.49 then traces exactly to the negative angular defect carried by each non-planar hexagon's own fan centroid — an auxiliary point the convention introduces but had not, until that subsection, been counted — restoring the exact total once every such point is (sections 25.7.48 and 25.7.49 and proposition 25.63); followed by the closing taxonomy, now on four axes rather than three, chapter 26) are likewise drafted, with a specific, itemized list of remaining open items — primary-source verification for every historical case, most recently Fuller and Sadao's own 1954 documentation and Peirce's own 1879 paper, both still unobtained; Gray's 1994/1995 transformation-equation papers, narrowed rather than closed by chapter 21's own reconstruction of the 1995 paper's specific construction from Gray's posted source code (section 21.5) and by appendix A's independent corroboration of that re-

# Contents

**Preface and Status of This Volume** **ii**

**I Foundations** **1**

**1 Introduction: What This Volume Is About** **2**
1.1 Three kinds of cartographic claim . . . . . 2
1.2 What this volume actually contains . . . . . 3
1.2.1 The metric core: curvature forces distortion (Chapters 2–5) . . . . . 3
1.2.2 The topological turn (Chapter 6), and its consequences for retroazimuthal projections (Chapter 7 and its addendum) . . . . . 3
1.2.3 How far the metric impossibility extends, and where it stops (Chapters 7b–7g) . . . . . 3
1.2.4 Named projections as worked instances of the impossibility (Chapters 8–10) . . . . . 4
1.2.5 Five case studies, each closing one cell of the eventual taxonomy . . . 4
1.2.6 Closing the bridge between the discrete and the continuous (Chapter 11d) . . . . . 4
1.2.7 The closing taxonomy, and a second verification pass (Chapter 11 and Appendix A) . . . . . 5
1.3 Why the metric and topological obstructions belong in one volume . . . . . 5
1.4 What is, and is not, claimed as original here . . . . . 6

**2 The Theorema Egregium: Curvature as an Intrinsic Obstruction** **7**
2.1 Historical note . . . . . 7
2.2 The first fundamental form . . . . . 7
2.3 The second fundamental form and extrinsic curvature . . . . . 8
2.4 Statement of the theorem . . . . . 8
2.4.1 Derivation via the Codazzi–Mainardi route . . . . . 8
2.4.2 A more usable closed form: the Brioschi formula . . . . . 9
2.5 Application to the sphere . . . . . 9

**3 Tissot's Indicatrix: Measuring the Unavoidable** **12**
3.1 Historical note . . . . . 12
3.2 Setup: the local linearization of a projection . . . . . 13
3.3 Principal directions and principal scales . . . . . 14
3.4 The three canonical distortion measures, read off the indicatrix . . . . . 15
3.5 The indicatrix as a manifestation of non-zero curvature . . . . . 15

**II The Impossibility Theorems Proper 17**

**4 No Map Is Both Conformal and Equal-Area 18**
4.1 Statement, and why it needs a chapter of its own . . . 18
4.2 The trade-off curve at a point . . . 19
4.3 From a point to a region: why the passage is not vacuous . . . 20
4.4 Taxonomic consequences: three families, no fourth . . . 20
4.5 Closing the conic case, and a correction proposition 4.6 itself needs . . . 22

**5 From Pointwise to Global: A Gauss–Bonnet Bound, and Milnor's Theorem 25**
5.1 What a "global" distortion bound should even mean . . . 25
5.2 A genuine bound: total curvature controls total conformal defect . . . 26
5.3 From flux to pointwise: a maximum-principle bound . . . 28
5.4 The quantitative bound: an exact solution on a geodesic disk . . . 29
5.5 The general boundary-value problem: an explicit Poisson kernel . . . 31
5.6 Milnor's theorem . . . 33
5.7 Resolving item (d): Milnor's paper contains two theorems, not one . . . 35
5.8 A closer reading of the Papadopoulos survey: Milnor's Chebyshev theorem in full, and a historical aside in Milnor's own words . . . 36
5.9 How the two results of this chapter relate . . . 38
5.10 Closing the gap: $\sigma_1, \sigma_2$ coincide with the pointwise extremes, on a geodesically convex region . . . 39

**6 The Hairy Ball Theorem: A Topological, Not Metric, Obstruction 42**
6.1 A different kind of impossibility . . . 42
6.2 Vector fields on a surface . . . 42
6.2.1 Proof via degree theory . . . 43
6.3 Application: why a global retroazimuthal direction field cannot exist . . . 44
6.4 What remains for item (iv) . . . 45

**7 Item (iv), Revisited: Conformality and Retroazimuthality Are Compatible, and the Only Obstruction Is Topological 46**
7.1 What item (iv) originally claimed . . . 46
7.2 What the companion volume actually establishes . . . 46
7.2.1 Chapter 13: the retroazimuthal condition as a Cauchy–Riemann system 47
7.2.2 Chapters 12–13: the local obstruction does not survive globally, and a closed form exists . . . 47
7.3 The domain restriction is forced by topology, not by a separate metric obstruction . . . 48
7.4 What remains open . . . 48
7.5 Uniqueness beyond the vector-alignment family, closed . . . 48
7.6 Revising the volume's four-item programme . . . 50

**8 The Equidistance Theorem: How Much Distance Can Ever Be Preserved 51**
8.1 Why this theorem has no chapter yet . . . 51
8.2 The impossibility of full equidistance . . . 51
8.3 Distortion-free loci and Tissot's indicatrix . . . 52
8.4 How many points can be simultaneously exact . . . 53
8.4.1 Two points: achievable . . . 53
8.4.2 Three points: generically impossible . . . 53
8.5 Where this leaves the taxonomy . . . 54

**9 Beyond the Sphere: Why the Impossibility Is the Sphere's, Not Every Surface's** **55**
9.1 A question this volume has so far begged . . . 55
9.2 The Gauss–Bonnet theorem for closed surfaces . . . 55
9.3 What Gauss–Bonnet forces at each genus . . . 57
9.4 Why this does not touch cartography, and why it should still be in this volume 57

**10 The Negative Case, and Why Conformality Never Needed the Sphere at All** **59**
10.1 Completing a trichotomy, cautiously . . . 59
10.2 The hyperbolic plane . . . 59
10.3 What genus $g \geq 2$ gets, that the sphere and torus do not . . . 60
10.4 Why conformal maps exist regardless, and what this says about the whole volume . . . 60

**11 Hilbert's Theorem: The Hyperbolic Plane Cannot Even Be Built** **62**
11.1 A question chapter 10 left unasked . . . 62
11.2 Beltrami's pseudosphere: a local realization, and where it fails . . . 63
11.3 The asymptotic net, and the sine-Gordon equation . . . 63
11.4 Hazzidakis's area formula . . . 65
11.5 Hilbert's theorem . . . 66
11.6 A second theorem in the same paper: rigidity of the sphere . . . 68
11.7 What this adds to the volume's picture of curvature sign . . . 69

**12 Efimov's Theorem: When Constant Curvature Was the Easy Case** **72**
12.1 Cohn-Vossen's question . . . 72
12.2 Efimov's theorem . . . 73
12.3 Why constant curvature was the easy case . . . 74
12.4 Efimov's argument, via the Gauss map . . . 75
12.5 The hypothesis is sharp: the catenoid . . . 76
12.6 A more dramatic failure of sharpness: Nadirashvili's surface . . . 78
12.7 What is possible in higher dimensions . . . 79
12.8 Later extensions, briefly . . . 80
12.9 What this adds to the volume's picture . . . 81

**13 The Nash–Kuiper Theorem: Regularity as the Loophole** **83**
13.1 An open thread from two earlier chapters . . . 83
13.2 Short maps, and the shape of the theorem . . . 84
13.3 Why this does not contradict Hilbert or Efimov . . . 84
13.4 The mechanism in miniature: corrugating a line segment . . . 85
13.5 Nash's own construction: twisting deformations . . . 86
13.6 Kuiper's refinement: stretching corrugations replace twists . . . 88
13.6.1 What reading the translations, rather than the originals, licenses . . . 89
13.7 The flat torus, corrugated explicitly . . . 90
13.8 Closing claim (a): no compact $C^2$ surface is flat . . . 91
13.9 What this adds to the volume's picture . . . 92

**III Historical Case Studies 93**

**14 Mercator and Lambert: Two Solutions, and an Eighteenth-Century Sighting of the Theorem 94**
14.1 Why Part III opens with this pair, and not with something else . . . . . 94
14.2 Mercator, 1569: solving conformality without a theory of conformality . . . . 94
14.3 Legend 3 of the 1569 map: an English translation of Mercator's own stated method, reached at one remove . . . . . 96
14.4 Lambert, 1772: the equal-area solution, and a qualitative statement of theorem 4.1 itself . . . . . 96
14.5 The primary text read directly, and a narrowing rather than a confirmation . 98

**15 Loxodromes: The Curve Mercator Was Actually Built to Straighten 100**
15.1 A complement, not a gap . . . . . 100
15.2 Definition and the differential equation . . . . . 100
15.3 The loxodromic spiral and its infinite length . . . . . 101

**16 Gauss, Krüger, and the Regional Strategy: Trading Global Coverage for Local Fidelity 102**
16.1 A different strategy, not a third family . . . . . 102
16.2 From Lambert's spherical sketch to Gauss's ellipsoidal survey mathematics . 102
16.3 Krüger's own monograph: an open-access digitization located, but not yet readable . . . . . 104
16.4 The central meridian as proposition 4.6 made concrete . . . . . 104
16.5 Placement and what Part III still owes . . . . . 105

**17 The Ellipsoidal Correction: Gauss–Krüger as It Is Actually Used 107**
17.1 What the sphere-based chapter left out . . . . . 107
17.2 The ellipsoid of revolution . . . . . 107
17.3 Isometric latitude and the conformal mapping . . . . . 108
17.4 Meridian arc length: the one place ellipsoidal geodesy has no closed form at all 110

**18 Discrete Global Grid Systems: Abandoning the Continuous Chart Altogether 111**
18.1 A third response, not a third family . . . . . 111
18.2 The founding survey: five design choices, one representative construction . . . 111
18.3 Twelve unavoidable defects: a discrete echo of chapter 6 . . . . . 112
18.4 Placement, and the close of Part III . . . . . 114

**19 Robinson, 1963: The Aphylactic Family Gets Its Case Study 115**
19.1 The gap this chapter closes, and the one it does not . . . . . 115
19.2 A projection built by eye, then fitted with numbers . . . . . 115
19.3 Confirming the family placement by direct computation . . . . . 117
19.4 Off the central meridian: distortion as a function of $\lambda$, not only $\varphi$ . . . . . 119
19.5 An orthogonal question: what interruption would, and would not, change . . 120
19.6 Placement, and where this leaves Part III . . . . . 121

**20 Goode, 1923: Interruption Gets Its Own Case Study, and Why the Seam Sits Where It Does 123**
20.1 Closing the second gap, with a case from a different family than the first . . . 123
20.2 The equatorial piece: an exact, closed-form proof of equivalence . . . . . 124

20.3 The polar piece: Mollweide, proved rather than cited . . . . . 124
20.4 Where the seam sits, and why: an exact join condition, verified numerically . 125
20.5 Family and continuity, together: what this case adds to the taxonomy . . . . 127

**21 Fuller and Sadao, 1943/1954: Closing the Aphylactic-and-Interrupted Cell 128**
21.1 The third gap, and why it is examined last . . . . . 128
21.2 History: from cuboctahedron to icosahedron . . . . . 128
21.3 This volume's construction: gnomonic projection onto the icosahedron's faces 130
21.4 Distortion: an exact, closed-form proof of aphylaxis . . . . . 131
21.5 Gray's exact equations, reconstructed and computed . . . . . 132
21.6 A denser distortion field, and a discrepancy the reconstruction itself surfaces . 133
21.7 Interruption, by construction . . . . . 135
21.8 The icosahedron's twelve vertices . . . . . 136
21.9 Family and continuity, together: the fourth cell . . . . . 136
21.10 Item (iii) attempted: a real-coastline check of the icosahedron this chapter's orientation fixes . . . . . 138

**22 Peirce, 1879: The Conformal Case, and Why It Resists the Fourth Cell 140**
22.1 An eighth case, examined for a different reason than the first seven . . . . . 140
22.2 History: an elliptic function for a compass point . . . . . 140
22.3 Conformality by composition: what this chapter proves, and what it cites . . 142
22.4 The Schwarz–Christoffel exponent rule, derived rather than cited . . . . . 144
22.5 The normalization constant, fixed . . . . . 147
22.6 Peirce's own two tables: a further access pass, narrowed but not closed . . . . 149
22.7 Peirce's own two tables, read: the scale question closed . . . . . 151
22.8 The four singular points, and their place in this volume's cone-point framework 153
22.9 The branch exponent, computed . . . . . 154
22.10 Conformal and contiguous . . . . . 155
22.11 The primary source, obtained . . . . . 156
22.12 Christoffel 1867 and Schwarz 1869, at last located: two different access outcomes for the same citation . . . . . 157
22.13 The 1877 Coast Survey report, located: the appendix is two pages, not the volume as a whole . . . . . 162
22.13.1 A further pass (6 Aug 2026): the oarcloud mirror reached, and a new, more precise diagnosis than a robots block . . . . . 164

**23 This Volume's Own Answer: An Interrupted Construction, Conformal on Every Gore 166**
23.1 Returning to the question chapter 22 left open . . . . . 166
23.2 The construction: stereographic projection, applied gore by gore . . . . . 166
23.3 Interruption, exactly as before . . . . . 167
23.4 What this closes, and what it does not . . . . . 168

**24 Three More Cases: A National Geographic Lineage, and a Controversy About What "Best" Means 169**
24.1 Why these three, together . . . . . 169
24.2 Van der Grinten (1898/1904): the circle that came before Robinson . . . . . 170
24.3 Winkel Tripel (1921): an average of two other projections . . . . . 171
24.4 Gall (1855) and Peters (1967/1973): one projection, two names, and a genuine equal-area proof . . . . . 173
24.4.1 Two names, and how the second one was affixed . . . . . 173

24.4.2 The controversy, presented on its own terms . . . 174
24.5 What this adds to the taxonomy, and what it does not . . . 175

**IV Synthesis 177**

**25 Descartes, Poincaré–Hopf, and the Twelve Defects: Closing the Bridge 178**
25.1 What section 18.3 left open . . . 178
25.2 Cone points on a polyhedral surface . . . 178
25.3 Holonomy around a cone point . . . 179
25.4 From holonomy to index: the cited bridge . . . 179
25.5 The twelve defects, re-derived from Poincaré–Hopf . . . 182
25.6 The bridge, stated formally . . . 183
25.7 A refining sequence: what converges immediately, and what does not . . . 183
25.7.1 The sixty defects, actually computed . . . 185
25.7.2 A face-flattening convention that restores exactness, and what it shows 187
25.7.3 The pentagon-corner decay rate, extended to $n = 20$: consistent with an inverse-square law, not derived as one . . . 188
25.7.4 A least-squares fit, extended to $n = 40$: closer to $-2$, still not equal to it 189
25.7.5 Extending the fit to $n = 50$: the confidence interval keeps tightening around $-2$, without reaching it . . . 190
25.7.6 Does the curvature distribution flatten out, or converge to a fixed shape? 191
25.7.7 An exact law for the mean defect, and what it explains . . . 193
25.7.8 Is $\kappa_\infty$ scheme-invariant? A three-way construction test . . . 195
25.7.9 Isolating the SSB/RMS/GEO gap: flat chords or nested normalization? 197
25.7.10 Why RMS and GEO still differ: nested slerp does not commute . . . 197
25.7.11 Propagating the mechanism: why GEO's $\kappa_\infty$ is not one local number . 199
25.7.12 A genuinely symmetric construction, built and tested: it does not land on RMS . . . 201
25.7.13 Sharpening the coincidence: it already holds at the level of the raw displacement, not just the extrapolated $\kappa_\infty$ . . . 202
25.7.14 The mechanism isolated directly: privilege/nearness mismatch, not vertex identity, drives the split . . . 203
25.7.15 High-precision values, and a genuine parity asymmetry between the two classes . . . 205
25.7.16 The parity puzzle, reduced: one exact lemma, one exact symmetry, one honest unknown . . . 206
25.7.17 The parity, proved: a 180° rotation, not an expansion . . . 207
25.7.18 Taking the reflection trick to RMS: two identities transplant exactly, the parity does not . . . 209
25.7.19 Closing the loop: does the role-mixture model actually predict the $\sim 3\%$ global gap? . . . 212
25.7.20 Toward a closed form for $\kappa_\infty$: a tighter numerical bound, and a negative screen . . . 214
25.7.21 The right existing machinery for this question, located rather than built 216
25.7.22 A further attempt to obtain the paper itself: bibliographic record sharpened, full text still not reached . . . 216
25.7.23 The reprint volume's own publisher record, reached directly: still no chapter-level text . . . 216

25.7.24 A candidate mechanism, tried directly: does the defect simply track a vertex-associated area? . . . 217
25.7.25 A second candidate, tried directly: does local shape explain the residual? 218
25.7.26 A third candidate, tried directly: is the residual an artifact of measuring area on the flat, rather than the curved, faces? . . . 220
25.7.27 Extending to $n = 100$: the model gap narrows, and $\kappa_\infty$ tightens to five figures . . . 221
25.7.28 A practical ceiling: $n = 150, 200$, and $\kappa_\infty$ to five figures . . . 223
25.7.29 Past the practical ceiling: a localized, $O(1)$-per-$n$ reconstruction reaches $n = 200{,}000$ . . . 224
25.7.30 The named next step, taken: arbitrary-precision arithmetic past the point-degeneracy ceiling . . . 226
25.7.31 The shape statistics, extended to $n = 100$: the plateau holds . . . 228
25.7.32 The shape statistics, extended once more: to $n = 120, 150, 200$ . . . 229
25.7.33 A first attempt at a convergence rate for $\mathrm{CV}(n)$ and $M(n)$, and why it does not yet succeed . . . 230
25.7.34 The postponed recomputation, completed: a vectorized rewrite reaches $n = 400$, and the convergence rate is no longer a guess . . . 231
25.7.35 The direct argument, found: an exact closed form for $\kappa_{\mathrm{other}}(n)$, not merely its asymptotics . . . 232
25.7.36 The second moment: no new identity is needed, because one was already in hand . . . 234
25.7.37 Why $\kappa_\infty^{\mathrm{SSB}}$'s localization trick cannot be transplanted to $M(n)$: the maximizing vertex is not corner-local at all . . . 235
25.7.38 A closed-form candidate for $M_\infty$: the flat-chord area Jacobian, computed exactly . . . 236
25.7.39 A closed-form candidate for $\mathrm{CV}_\infty$ too: no argmax is needed for a variance 239
25.7.40 The centroid locus, localized after all: an $O(1)$-per-$n$ construction reaches $n = 10{,}000$, and a new, differently-caused ceiling . . . 240
25.7.41 The fixed-triangle evaluation, in arbitrary precision: the float64 floor removed, not merely lowered . . . 242
25.7.42 The coefficient in closed form: $c = 405\sqrt{3}(47 - 21\sqrt{5})/(4\pi)$, to 69 stable digits . . . 245
25.7.43 The search itself, in arbitrary precision: the selection-bias hypothesis tested directly . . . 247
25.7.44 Does the same mechanism reach $\kappa_\infty^{\mathrm{SSB}}$? A structural answer, not a new computation . . . 248
25.7.45 The fixed-point push, taken up: a real bug found on the way, and $\kappa_\infty^{\mathrm{SSB}}$ pushed to roughly fifty stable digits . . . 249
25.7.46 A second basis, from the golden ratio and the icosahedral central angle directly . . . 251
25.7.47 A third basis, from the flat-chord construction's own normalization step 252
25.7.48 A third convention: centroid-fan triangulation . . . 253
25.7.49 Closing the gap: the face centroids, counted as their own cone points . 254

**26 The Closing Taxonomy: Four Axes, One Proved, Three Organizing 258**
26.1 What this chapter can, and cannot, claim . . . 258
26.2 The full grid, populated . . . 259
26.3 The edge case that motivated the fourth axis . . . 260
26.4 How the taxonomy relates to this volume's four impossibility theorems . . . 261

26.5 A fifth axis proposed, and four chapters outside the grid entirely . . . . . 262
26.6 The fifth axis, closed: the exact-scale locus of every row in the table . . . . . 263
26.6.1 Two loci, not one, and why the distinction matters . . . . . 263
26.6.2 The nine rows . . . . . 263
26.6.3 Summary . . . . . 268
26.7 A discrete analogue for the ISEA3H family: do any cell edges match their geodesic arc? . . . . . 269
26.8 Snyder's actual curved boundaries: which, if any, are geodesics? . . . . . 270
26.9 Five more cells, populated from outside this volume's own case studies . . . . 273
26.10 Status of the volume at the close of this draft . . . . . 276

**A A Second Verification Pass** **283**
A.1 Purpose of this appendix . . . . . 283
A.2 Milnor, 1969 . . . . . 283
A.3 Ipbüker, 2004 . . . . . 285
A.4 Sahr, White, and Kimerling, 2003: the full body, obtained and read . . . . . 286
A.5 Peirce (1879) and Schwarz–Christoffel: an access path found, two corrections precisely dated . . . . . 287
A.6 White, Kimerling, Sahr, and Song, 1998: precisely identified, still paywalled, and its abstract mined for what it actually claims . . . . . 289
A.7 What this pass changes, and what it does not . . . . . 290
A.8 Two items added by the five later chapters . . . . . 292
A.8.1 A pass on Close (1921): access route confirmed, article text not yet extracted . . . . . 292
A.8.2 Maurer (1919): a plausible open access route identified, record not located . . . . . 293
A.8.3 A further pass on Maurer (1919): the catalog's own record numbering resists a shortcut, and the platform's public status is itself now less certain than assumed . . . . . 293
A.8.4 A pass on Poincaré (1907) and Koebe (1907): precise citations, one open discrepancy, access blocked mid-pass . . . . . 294
A.8.5 The 63-versus-64 discrepancy resolved, from a third independent source 294
A.9 Two items added by chapter 11 . . . . . 295
A.10 Three items added by chapter 13 . . . . . 297
A.11 One item added by chapter 12 . . . . . 298

# Part I

# Foundations

# Chapter 1

# Introduction: What This Volume Is About

## 1.1 Three kinds of cartographic claim

Cartographic writing — including several of the companion volumes in this series — tends to make claims of three different logical kinds, and it is worth separating them before going further, because only one of the three is the subject of this book.

The first kind is *empirical*: a specific projection, applied to a specific region, produces so much areal distortion, or so much angular distortion, at such-and-such a location. This is the kind of claim made throughout *No Map Is Best* and the DGGS monograph, and it is verified by direct computation — the finite-difference Tissot analysis, the Kerkovits band comparisons — against a specific numerical input.

The second kind is *historical*: a named projection was devised at a named date, for a named purpose, by a named author, and its adoption or rejection followed a documented institutional path. This is the register of *The Retroazimuthal Problem*'s treatment of the qibla tradition and of *The New Datum*'s account of the NSRS modernization. This volume's own case-study chapters (section 1.2.5 below) adopt this register too, wherever the historical record settles a question the mathematics alone leaves open.

The third kind is *impossibility*: no projection of a given class can have a given property, full stop, independent of cleverness, computing power, or future discovery. This is a much stronger and much rarer kind of claim, because it is a universal negative over an infinite class of maps, and a universal negative can only be secured by a proof, not by an exhaustive search. This volume is built around claims of the third kind, together with two further things a catalogue of impossibility results demands once it is taken seriously: an honest account of exactly where each impossibility's boundary lies (a regularity class relaxed, a domain restricted, a metric traded for a topological invariant), and, where the boundary itself turns out to be interesting, an explicit construction living right at it.

## 1.2 What this volume actually contains

**Epistemic Notice**

This section replaces an earlier draft's four-item table of contents (items (i)–(iv), all marked "forthcoming" at that stage) with an accurate account of the volume as it now stands. Nothing in the four-item plan was abandoned — every one of the four results it named is proved somewhere below — but the volume grew substantially past that original scope as later chapters kept finding that closing one impossibility cleanly required either a genuinely different mathematical tool (topology rather than curvature, for item (iii); a positive construction rather than a further negative result, for the Nash–Kuiper material) or a concrete worked example anchoring the abstract classification (the five case-study chapters, section 1.2.5). The volume's own discipline, stated in section 1.4 below and enforced throughout by the `todobox/disclaimerbox` convention used in every chapter, is to report what has actually been proved, verified, or found in primary sources, and to flag explicitly — not silently smooth over — whatever remains open, unverified, or contradicted by later work.

### 1.2.1 The metric core: curvature forces distortion (Chapters 2–5)

Chapter 2 establishes the foundational fact every later impossibility in this volume ultimately rests on: Gauss's Theorema Egregium, that the Gaussian curvature of a surface is an intrinsic invariant, unchanged by any length-preserving deformation, so that a sphere (curvature $1/R^2 \neq 0$ everywhere) cannot be isometric to any flat region of the plane. Chapter 3 gives the local measuring instrument built on top of that fact — Tissot's indicatrix, the ellipse into which an infinitesimal circle on the sphere is carried by a given projection — from which every scalar distortion measure used elsewhere in this series can be read off directly. Chapter 4 proves the first genuine impossibility result, that no map preserves both angles and areas simultaneously (short of a degenerate region), and chapter 5 turns Tissot's local measure into a global integral inequality, Milnor's sharp lower bound on total distortion, closing the volume's original items (i) and (ii).

### 1.2.2 The topological turn (Chapter 6), and its consequences for retroazimuthal projections (Chapter 7 and its addendum)

Chapter 6 shifts the obstruction from metric to topological: the hairy ball theorem, that no continuous nowhere-zero tangent vector field can exist on the whole sphere, rules out a naive global construction for retroazimuthal projections that works perfectly locally. Chapter 7 and its addendum (section 7.5, via `chapters/ch07_addendum_retro_uniqueness.tex`) then settle the compatibility question the original item (iv) posed: conformality and retroazimuthality *are* jointly achievable, for an arbitrarily prescribed centre, on any domain no larger than the sphere minus that centre and its antipode, and the only genuine obstruction turns out to be exactly chapter 6's topological one, not a further metric one. This is a compatibility result, not a further impossibility, and revising the original plan to say so explicitly — rather than forcing a fourth negative result where the mathematics does not support one — is itself an instance of the reporting discipline section 1.4 describes.

### 1.2.3 How far the metric impossibility extends, and where it stops (Chapters 7b–7g)

A cluster of shorter chapters interrogates chapters 4 and 5's metric impossibility from every direction a careful reader might press on it, rather than leaving it as a single isolated theorem. Chapter 8 asks how much distance, rather than angle or area, can ever be preserved, and proves a companion bound. Chapter 9 asks whether the impossibility is a fact about

the sphere specifically or about curved surfaces generally, and traces it to genus via Gauss–Bonnet. Chapters 10 to 12 take up the negative-curvature case in turn — first showing conformal flattening never actually needed positive curvature at all, then Hilbert's theorem that the full hyperbolic plane cannot even be isometrically embedded in Euclidean 3-space, then Efimov's sharper theorem identifying exactly which negatively-curved surfaces resist embedding and which do not, the constant-curvature case turning out to be the easy one rather than the hard one. Chapter 13 closes this cluster with the volume's one genuine counternote within the impossibility theme: the Nash–Kuiper theorem, that dropping smoothness down to merely $C^1$ regularity turns an impossible isometric embedding into an achievable (if wildly non-smooth, "corrugated") one — a loophole in the regularity class rather than a refutation of anything proved earlier, and included precisely because a volume about impossibility results owes its reader the one prominent case where relaxing a single hypothesis flips the answer.

### 1.2.4 Named projections as worked instances of the impossibility (Chapters 8–10)

Chapter 14 and its companion on loxodromes (chapter 15) ground the abstract conformal/equal-area trade-off in the two most historically consequential single-purpose solutions to it, including an eighteenth-century anticipation of the impossibility theorem itself found in the primary-source record. Chapters 16 and 17 take up the regional strategy — trading global coverage for local fidelity — first on the sphere and then, as the projection is actually used in practice, on the ellipsoid. Chapter 18 pushes the trade-off to its logical extreme, the discrete global grid systems that abandon the continuous chart altogether, connecting this volume's impossibility results to the same H3/S2/rHEALPix/DGGRID material developed at length in the companion DGGS monograph, *From Continuous to Discrete*.

### 1.2.5 Five case studies, each closing one cell of the eventual taxonomy

Five chapters (19 to 23) take up specific named projections not as further abstract instances but as historically and mathematically concrete case studies, each one closing a specific cell of chapter 26's classification: Robinson (1963) for the deliberately-aphylactic compromise family; Goode (1923) for interruption, including why the seam sits where it does rather than somewhere else equally plausible; Fuller and Sadao (1943/1954) for the aphylactic-and-interrupted combination; Peirce (1879) for the conformal case, and for why it resists being forced into a fourth, different cell the way the first three were; and, finally, this volume's own contribution rather than a historical case study — an interrupted construction, conformal on every individual gore, proposed in chapter 23 as this volume's own answer to the question the Peirce chapter leaves standing. Chapter 22's own addenda (`ch11e_addendum_schwarz_christoffel_derivation.tex`, deriving the Schwarz–Christoffel exponent rule from scratch rather than citing it, and `ch11e_addendum_peirce_normalization_constant.tex`) are where the case study's mathematical apparatus, rather than its historical narrative, is actually built out.

### 1.2.6 Closing the bridge between the discrete and the continuous (Chapter 11d)

Chapter 25 is the volume's most extensive single chapter, and the one most recently under active development: it establishes, via Descartes's angle-defect theorem and the Poincaré–Hopf index theorem, a formal (not merely observed) connection between the discrete cur-

vature carried by the twelve icosahedral disclinations of a subdivided Goldberg polyhedron — exactly the construction chapter 18 uses for one class of discrete global grid — and the continuous hairy-ball obstruction of chapter 6. Its own internal discipline is the strictest in the volume: every numerical claim is computed by an independently-checked script under `scripts/`, every convention (raw-angle, plane-projected, centroid-fan) is tested against the others rather than assumed, and open questions are carried forward explicitly across many chapter drafts rather than quietly dropped. Its most recent addition (section 25.7.30) replaced a float64 numerical ceiling on the pentagon-corner curvature ratio $\kappa_\infty^{\mathrm{SSB}}$ with an arbitrary-precision reconstruction reaching $n = 10^{10}$, tightening that constant from five significant figures to eighteen self-consistent ones ($\kappa_\infty = 0.602823502033767116\ldots$) while leaving its closed-form identity — which a wide automated search, including `mpmath`'s PSLQ-based `identify` routine, still finds no match for — honestly open, exactly the kind of numerically-tightened-but-not-closed result this volume's discipline asks for rather than a claimed derivation it cannot support. Chapter 25's own companion addendum, `ch11d_addendum_pentagon_corner_jacobian.tex`, carries out a related closed-form computation (for a different quantity, $M_\infty$, where the area-Jacobian argument does close exactly) that the main chapter explicitly contrasts with $\kappa_\infty^{\mathrm{SSB}}$'s own, still-unclosed case.

### 1.2.7 The closing taxonomy, and a second verification pass (Chapter 11 and Appendix A)

Chapter 26 is the volume's intended synthesis: a classification of named projections along four axes (conformal/equal-area/aphylactic; global/regional; injective/interrupted; and a fourth, this volume's own contribution rather than a citation from the literature), with one axis — the trade-off proved in chapters 4 and 5 — resting on a proved theorem and the remaining three organizing the historical record without themselves being further impossibility results. Its own addendum, section 26.6 (`ch11_taxonomy_addendum_fifth_axis.tex`), closes a fifth axis — the exact-scale locus of every row in the table — left open in the chapter's first draft. Section 26.10, within chapter 26, is this volume's running ledger of exactly what remains open across every chapter; a reader wanting a single, current, honest inventory of unfinished business should consult that section directly rather than this introduction, which is necessarily a snapshot and will drift out of date again as work continues. Finally, appendix A is a second, independent verification pass over the volume's own computed results, run in the same spirit chapter 25 already applies internally — checking published numbers against fresh, independently-written scripts rather than only against themselves.

## 1.3 Why the metric and topological obstructions belong in one volume

The four original impossibility results — and every later chapter that extends, tests, or grounds them historically — are not simply thematically related; they share a single underlying mechanism, which this volume tries to make visible rather than leaving it implicit the way most cartographic textbooks do. In every impossibility proved below, the obstruction is a fact about the *intrinsic geometry of the sphere* — its Gaussian curvature, in the metric results; its topology, in the hairy-ball and defect-bridge results — that no choice of projection formula can work around, because a projection is by definition a correspondence between points, and it cannot alter the metric or topological facts that hold prior to and independent of any such correspondence. This is the reason a flat map of any part of the Earth's surface must misrepresent length, angle, or area to some degree, however small the region, and the reason no retroazimuthal construction can be built by a naive vector-field

argument valid everywhere at once. The case-study chapters and the closing taxonomy exist because an impossibility theorem, stated only in the abstract, risks reading as a curiosity; showing exactly which historical projections chose which side of each trade-off, and why, is what turns the abstract results into an account of the actual history of the discipline.

## 1.4 What is, and is not, claimed as original here

The core theorems in this volume — Gauss's Theorema Egregium, Tissot's indicatrix, Milnor's distortion bound, the hairy ball theorem, Hilbert's and Efimov's embedding theorems, the Nash–Kuiper theorem — are all more than fifty years old and none of them is new. What this volume attempts, in keeping with the working method of the rest of the series, is: (a) to derive each one explicitly rather than citing it as a black box, so that the connection between the abstract proof and the concrete cartographic phenomenon is traceable step by step — carried furthest in chapters 22 and 25's addenda, which derive results such as the Schwarz–Christoffel exponent rule and the pentagon-corner curvature machinery from scratch rather than citing them; (b) to state precisely which projections and constructions in the companion volumes each theorem forecloses, rather than leaving the connection at the level of loose analogy; (c) to be explicit, at each step, about what has been proved, what has been verified numerically to a stated precision, and what has merely been made plausible — a distinction this series has tried to maintain since the first volume of *No Map Is Best*; and (d), where a numerical or historical claim genuinely cannot be settled within a given pass — a primary source not reached, a closed form not found despite a real search — to record that fact honestly, with what was tried and why it did not close the question, rather than either overstating the result or silently omitting the attempt. Chapter 25's long sequence of `todobox` entries, several carried open across many chapter drafts before being narrowed or closed by a later one, is this discipline in its most sustained form anywhere in the volume.

# Chapter 2

# The Theorema Egregium: Curvature as an Intrinsic Obstruction

## 2.1 Historical note

Gauss presented his *Disquisitiones generales circa superficies curvas* to the Royal Society of Göttingen on 8 October 1827; it was published the following year in the *Commentationes Societatis Regiae Scientiarum Gottingensis*, volume 6, pp. 99–146.[1] The memoir grew directly out of Gauss's work, from 1821 to 1825, as director of the geodetic survey of the Kingdom of Hanover — a fact this series has occasion to note also in *The New Datum*'s account of geodetic reference frames, since Gauss's own practical grappling with the discrepancy between ellipsoidal measurement and planar representation is the historical seed of the entire subject of this volume. The paper's central result is the one Gauss himself named the *theorema egregium*, the "outstanding" or "remarkable" theorem, and the name has stuck for very nearly two hundred years.

## 2.2 The first fundamental form

Let $S \subset \mathbb{R}^3$ be a smooth surface, parametrized locally by $\mathbf{x}(u, v)$ for $(u, v)$ in some open subset of $\mathbb{R}^2$. Write $\mathbf{x}_u = \partial \mathbf{x}/\partial u$ and $\mathbf{x}_v = \partial \mathbf{x}/\partial v$. The *first fundamental form* is the quadratic form on tangent vectors induced by the ambient Euclidean inner product:

$$I(a, b) = E\,a^2 + 2F\,ab + G\,b^2, \qquad E = \mathbf{x}_u \cdot \mathbf{x}_u, \quad F = \mathbf{x}_u \cdot \mathbf{x}_v, \quad G = \mathbf{x}_v \cdot \mathbf{x}_v.$$

This form computes the length of any tangent vector, and hence, by integration, the length of any curve drawn on $S$: it is the object that encodes everything about distances and angles *as measured along the surface*, with no reference to how $S$ happens to sit inside $\mathbb{R}^3$. Two surfaces (or two regions of the same surface, under different parametrizations) are called *isometric* if there is a correspondence between them under which $E$, $F$, $G$ agree at corresponding points for all curves; equivalently, an isometry is a map that preserves the length of every curve.

A map projection, in exactly the sense used throughout the rest of this series, *is* a choice of coordinates $(u, v)$ for a region of the sphere together with the decision to plot $(u, v)$ directly

[1] Consulted via the Internet Archive digitization of the original (archive.org, item `bub_gb_bXOAAAAAMAAJ`) and cross-checked against the historical accounts in Merker (2014) and the University of Padua seminar note on the two proofs Gauss gave of the invariance theorem, one in 1825 (published posthumously) and the definitive one in the 1827 memoir itself.

on flat paper. The projection is length-preserving (Gauss's technical term, still current, is that it produces a *developable* or *isometric* image) if and only if the first fundamental form of the sphere in those coordinates is literally the flat Euclidean form $du^2 + dv^2$ everywhere on the region in question. The question this chapter answers is: for which surfaces is this possible?

## 2.3 The second fundamental form and extrinsic curvature

The first fundamental form alone does not distinguish a flat sheet of paper from the same sheet rolled into a cylinder — both have $E = 1$, $F = 0$, $G = 1$ in the natural parametrization, and indeed a cylinder can be produced from a plane by bending without stretching or tearing, so the two are isometric. To capture how a surface curves *within* the ambient space, Gauss introduced the unit normal $\mathbf{n} = (\mathbf{x}_u \times \mathbf{x}_v)/|\mathbf{x}_u \times \mathbf{x}_v|$ and the *second fundamental form*

$$II(a, b) = e\,a^2 + 2f\,ab + g\,b^2, \qquad e = \mathbf{x}_{uu} \cdot \mathbf{n}, \quad f = \mathbf{x}_{uv} \cdot \mathbf{n}, \quad g = \mathbf{x}_{vv} \cdot \mathbf{n}.$$

The *Gaussian curvature* at a point is then defined, extrinsically, as

$$K = \frac{eg - f^2}{EG - F^2}. \tag{2.1}$$

On this definition alone, there is no reason to expect $K$ to be preserved under bending: $E, F, G$ depend only on the surface's internal geometry, but $e, f, g$ manifestly depend on the normal $\mathbf{n}$, i.e. on how the surface is embedded. A cylinder bent from a plane, for instance, has $e = f = g = 0$ trivially for the plane, but non-zero $e$ or $g$ once curved — and yet, as we are about to see, the *combination* $eg - f^2$ divided by $EG - F^2$ turns out to depend on $E, F, G$ and their derivatives alone.

## 2.4 Statement of the theorem

**Theorem 2.1** (Gauss, 1827)**.** *The Gaussian curvature $K$ defined by eq.* (2.1) *can be expressed entirely in terms of $E$, $F$, $G$ and their first and second partial derivatives with respect to $u$ and $v$. Consequently, $K$ is invariant under any isometry: if two surfaces (or two parametrizations of the same region) have the same first fundamental form, they have the same Gaussian curvature at corresponding points.*

### 2.4.1 Derivation via the Codazzi–Mainardi route

The cleanest modern route to theorem 2.1 — not Gauss's original argument, which proceeded by direct and formidable computation, but the equivalent route standard since Ricci and Levi-Civita, and adopted here because it makes the logic of the proof, rather than only its conclusion, transparent — runs through the Gauss and Codazzi–Mainardi equations for the derivatives of the frame $\{\mathbf{x}_u, \mathbf{x}_v, \mathbf{n}\}$.

Differentiating $\mathbf{x}_u, \mathbf{x}_v$ again and expressing the result in the (still tangent-plus-normal) frame gives the *Gauss equations*

$$\mathbf{x}_{uu} = \Gamma^1_{11}\mathbf{x}_u + \Gamma^2_{11}\mathbf{x}_v + e\,\mathbf{n}, \qquad \mathbf{x}_{uv} = \Gamma^1_{12}\mathbf{x}_u + \Gamma^2_{12}\mathbf{x}_v + f\,\mathbf{n}, \qquad \mathbf{x}_{vv} = \Gamma^1_{22}\mathbf{x}_u + \Gamma^2_{22}\mathbf{x}_v + g\,\mathbf{n},$$

where the $\Gamma^k_{ij}$ are the Christoffel symbols, determined entirely by $E, F, G$ and their first derivatives via the standard formulas

$$\Gamma^1_{11} = \frac{GE_u - 2FF_u + FE_v}{2(EG - F^2)}, \qquad \Gamma^2_{11} = \frac{2EF_u - EE_v - FE_u}{2(EG - F^2)},$$

and similarly for the remaining four symbols (omitted here for brevity; they follow the same pattern applied to $\Gamma^1_{12}, \Gamma^2_{12}, \Gamma^1_{22}, \Gamma^2_{22}$). The key structural fact is that *every* $\Gamma^k_{ij}$ is a function of $E, F, G$ and their $u$- and $v$-derivatives only — the second fundamental form does not enter.

Now impose the integrability (mixed-partials) condition $(\mathbf{x}_{uu})_v = (\mathbf{x}_{uv})_u$. Expanding both sides using the Gauss equations above, and using $\mathbf{n}_u, \mathbf{n}_v$ expressed via the Weingarten equations in terms of $e, f, g, E, F, G$, and collecting the coefficient of $\mathbf{x}_v$ on each side, one obtains (this is the classical computation; see Do Carmo, *Differential Geometry of Curves and Surfaces*, §4.3, for the full expansion) the relation

$$E\big(eg - f^2\big) = EG - F^2\,\left[\big(\Gamma^2_{11}\big)_v - \big(\Gamma^2_{12}\big)_u + \Gamma^1_{11}\Gamma^2_{12} - \Gamma^1_{12}\Gamma^2_{11} + \Gamma^2_{11}\Gamma^2_{22} - \Gamma^2_{12}\Gamma^2_{12}\right]. \tag{2.2}$$

The right-hand side, by construction, is built exclusively out of $E, F, G$ and their first and second $u, v$-derivatives (through the $\Gamma^k_{ij}$ and their derivatives). Dividing both sides by $E(EG - F^2)$ and comparing with eq. (2.1) gives $K$ expressed purely in terms of $E, F, G$: this is exactly theorem 2.1.

### 2.4.2 A more usable closed form: the Brioschi formula

For direct computation, it is more convenient to quote the equivalent closed-form determinant expression (Brioschi, 1852, an explicit repackaging of the content of eq. (2.2)):

$$K = \frac{\begin{vmatrix} -\frac{1}{2}E_{vv} + F_{uv} - \frac{1}{2}G_{uu} & \frac{1}{2}E_u & F_u - \frac{1}{2}E_v \\ F_v - \frac{1}{2}G_u & E & F \\ \frac{1}{2}G_v & F & G \end{vmatrix} - \begin{vmatrix} 0 & \frac{1}{2}E_v & \frac{1}{2}G_u \\ \frac{1}{2}E_v & E & F \\ \frac{1}{2}G_u & F & G \end{vmatrix}}{(EG - F^2)^2}. \tag{2.3}$$

For an *orthogonal* parametrization ($F \equiv 0$), which suffices for every case treated later in this volume and in the companion series, eq. (2.3) collapses to the much more tractable

$$K = -\frac{1}{2\sqrt{EG}}\left[\frac{\partial}{\partial u}\left(\frac{G_u}{\sqrt{EG}}\right) + \frac{\partial}{\partial v}\left(\frac{E_v}{\sqrt{EG}}\right)\right]. \tag{2.4}$$

**Key Result**

*Equation* (2.4) *is the form actually used, in practice, in Chapter 2 (Tissot indicatrix via finite differences) of the DGGS monograph and in the projection-by-projection curvature checks of* No Map Is Best*: given only the metric coefficients $E(u,v), G(u,v)$ of a projection's source surface in orthogonal coordinates, the curvature follows without any reference to how, or whether, that surface is embedded in space.*

## 2.5 Application to the sphere

Parametrize the sphere of radius $R$ by geographic coordinates $(\varphi, \lambda)$ (latitude, longitude), $\mathbf{x}(\varphi, \lambda) = R(\cos\varphi\cos\lambda,\ \cos\varphi\sin\lambda,\ \sin\varphi)$. A direct computation gives $E = R^2$, $F = 0$, $G = R^2\cos^2\varphi$. Substituting into eq. (2.4) with $u = \varphi$, $v = \lambda$:

$$\sqrt{EG} = R^2\cos\varphi, \qquad \frac{G_u}{\sqrt{EG}} = \frac{-2R^2\cos\varphi\sin\varphi}{R^2\cos\varphi} = -2\sin\varphi, \qquad \frac{E_v}{\sqrt{EG}} = 0,$$

$$K = -\frac{1}{2R^2\cos\varphi}\cdot\frac{\partial}{\partial\varphi}(-2\sin\varphi) = -\frac{1}{2R^2\cos\varphi}\cdot(-2\cos\varphi) = \frac{1}{R^2}.$$

The Gaussian curvature of a sphere of radius $R$ is the constant $1/R^2$, positive and strictly non-zero at every point, for any finite $R$. The plane, by the same formula applied to Cartesian coordinates ($E = G = 1$, $F = 0$, all derivatives zero), has $K \equiv 0$.

**Corollary 2.2** (No isometric flattening of the sphere)**.** *No non-empty open region of a sphere is isometric to any region of the plane. Consequently, no map projection — meaning, in the strict sense of chapter 1, a correspondence between a region of the sphere and a region of the plane — can preserve all distances simultaneously; some distortion, in length, angle, or area, is unavoidable for every projection of every non-degenerate region, however small.*

*Proof.* By theorem 2.1, $K$ is an isometric invariant. If a region $U$ of the sphere were isometric to a region $V$ of the plane, the invariant $K$ would have to agree at corresponding points: $1/R^2 = 0$, a contradiction, for any $R < \infty$. □

This corollary is the primitive fact underlying every subsequent chapter of this volume and, more broadly, every distortion computation performed anywhere else in this series. It is worth being precise about what it does and does not say. It does not say that distortion must be large; the curvature $1/R^2$ is a fixed, small number for a body the size of the Earth, and a sufficiently small region can be made to have arbitrarily small distortion by an appropriate projection (this is why local cadastral surveying can, in practice, often ignore curvature altogether, a point taken up at length in *The New Datum*'s treatment of SPCS2022 zone design). What corollary 2.2 rules out, absolutely and for regions of any size, is distortion being reduced all the way to *zero*. Chapters 3 and, in the planned continuation, 5 are about exactly how much distortion is forced, not merely that some is.

**Todo — verify primary source**

Confirm, against a facsimile or critical edition of the 1827/1828 Latin text itself (rather than only secondary historical accounts), the precise numbering and phrasing Gauss used for the theorem's statement in §12 of the *Disquisitiones*, before quoting or closely paraphrasing his own formulation in any later revision of this chapter.

**Key Result**

*[The article number and the theorem's own wording, confirmed directly] The numbering is Art. 12, not a different article: in the 1902 Princeton translation by Morehead and Hiltebeitel — prepared, as its translators state explicitly,* from a copy of the original Latin edition *(edition Ia, Gottingae, 1828) rather than from any later reprint — the theorem stands as the closing paragraph of the article numbered "12." in every Latin edition the translators collated (Ia, the Liouville reprint, and the Werke editions), immediately after the derivation, from the linear-element formula $\sqrt{E\,dp^2 + 2F\,dp\,dq + G\,dq^2}$ alone, of the fact that a surface's development leaves its curvature measure unchanged.*[a] *The Latin original — independently attested, in identical wording, by a scholarly citation to page 24 of the 1828 printing and by the standard secondary paraphrase this chapter had previously relied on — reads:* Si superficies curva in quamcumque aliam superficiem explicatur, mensura curvaturae in singulis punctis invariata manet.*[b] The English of Art. 12 renders this as: "If a curved surface is developed upon any other surface whatever, the measure of curvature in each point remains unchanged." This matches, word for word in substance, the paraphrase this chapter already gave of theorem 2.1's content, so nothing in the chapter's mathematical statement needs revision; what changes is only that the theorem's own words, and their location, are now sourced to a direct rendering of the primary Latin text rather than to a second-hand report of it.*

[a] Consulted directly via Project Gutenberg EBook #36856 (*General Investigations of Curved Surfaces of 1827 and 1825*, transcribed from the Boston University Library's Internet Archive scan of the 1902 Princeton edition), Art. 12, p. 20 of the translation.

[b] Cross-checked against a direct quotation, with page citation to the 1828 edition (p. 24), in a mathematical-physics monograph's discussion of the theorem, and against the Italian Wikipedia article on the *theorema egregium*, which reproduces the same Latin sentence; the two sources are independent of each other and of the Princeton translators, and agree letter for letter.

**Epistemic Notice**

Two caveats on section 2.5's sourcing remain worth stating. First, this chapter has not itself read the Latin of the 1828 *editio princeps* directly; it relies on a translation whose translators state they worked from a copy of that edition, corroborated by an independent Latin quotation with a page citation to the same edition, but a facsimile of the original has not been opened in this pass. Second, Gauss's own numbering convention in the 1827/1828 print is a bare Arabic numeral followed by a period ("12.") marking each new paragraph of the memoir, not a section sign; "§12", used informally in secondary literature and in this book's earlier todobox, is a modern convenience for referring to Art. 12, not Gauss's own notation, and later references in this volume should say "Art. 12" or "Article 12" when precision about the original matters.

# Chapter 3

# Tissot's Indicatrix: Measuring the Unavoidable

## 3.1 Historical note

Nicolas-Auguste Tissot (1824–1897) first presented the construction that bears his name in notes to the Académie des Sciences in 1859, but it was his full memoir, *Mémoire sur la représentation des surfaces et les projections des cartes géographiques*, serialized in the *Nouvelles annales de mathématiques* across 1878–1880 and published in consolidated book form by Gauthier-Villars in Paris in 1881, that established the indicatrix as the standard tool it remains today.[1] Tissot's own summary of his central result — that at every point of a surface being mapped, there exists at least one pair of directions, orthogonal on the source surface, that remains orthogonal in the image, however distorted the map is otherwise — is quoted, at second hand via the standard secondary literature (Snyder 1993; the Wikipedia synthesis of the primary claim), rather than from the primary 1881 text directly.

**Todo — verify primary source**

The characterization of Tissot's own statement of the orthogonal-directions theorem is currently sourced to secondary accounts (Snyder 1993 as excerpted in tertiary summaries) rather than to a direct reading of the 1881 *Mémoire*. Locate and consult a digitized copy of the primary text directly (the Internet Archive item above is a candidate) before treating the exact phrasing as settled.

[1] Publication details confirmed against the Archive for History of Exact Sciences study of Tissot by A'Campo, Gray and Laudenbach's collaborators (Springer, 2016), the Internet Archive digitization of the 1881 volume itself (`mmoiresurlarepr00tissgoog`), and the HathiTrust catalogue record for the same edition.

**Key Result**

*[A scholarly reading of the primary memoir, with page citations, replaces the tertiary paraphrase] Papadopoulos's 2016 study of Tissot — a mathematician's direct account of the 1881* Mémoire*, citing specific pages of the memoir itself (p. 1 for its opening remark on projection choice, p. 21 onward for the indicatrix's axis formulae, p. 71 for the projection later used in the Egyptian cadastral survey) rather than reporting Tissot only through Snyder's later paraphrase — gives the theorem in something close to Tissot's own words, carried over essentially unchanged from his first 1859* Comptes Rendus *note into the proved, expanded version that opens Chapter 1 of the 1881 memoir.*[a] *Papadopoulos's own translation of Tissot's numbered first principle reads:* "For any kind of representation of a surface onto another, there exists, at every point of the first surface, two perpendicular tangents, which are unique unless the angles are preserved at that point, such that the images of these two tangents are perpendicular on the second surface." *Tissot's own name for these two directions, per the same source, is* tangentes principales *(principal tangents) — the direct antecedent of this chapter's "principal directions" in theorem 3.1 — and he pairs the existence-and-uniqueness statement with a second numbered principle identifying those directions as exactly the ones along which the length-scaling ratio attains its extreme values $a, b$, matching theorem 3.1 in this chapter almost clause for clause. A companion, more general remark recorded from the same chapter — that varying the choice of orthogonal net gives infinitely many decompositions into infinitesimal rectangles, of which the principal-tangent pair is the unique one preserved as an orthogonal pair by the map itself — sharpens "unique unless angle-preserving" into the stronger and more precise statement that uniqueness holds for the map's own canonical net, not merely for some net.*

[a] A. Papadopoulos, *Nicolas-Auguste Tissot: A link between cartography and quasiconformal theory*, arXiv:1612.00279 (2016), §§2–3; the paper states explicitly that "the first chapter [of the 1881 memoir] contains, besides the proofs of the results announced in the *Comptes Rendus* notes, a description of the Tissot indicatrix" — i.e. the 1859 statement quoted below is the one the 1881 memoir goes on to prove, not a distinct earlier claim later abandoned.

**Epistemic Notice**

What section 3.1 closes, and what it leaves open, should be stated precisely. It closes the specific gap the todobox named: the theorem's substance and its attribution to Tissot's own numbered *principe* 1 are now sourced to a scholarly reading of the memoir itself, with page references, rather than to Snyder's tertiary paraphrase alone — and the wording quoted is close enough to this chapter's own theorem 3.1 that no revision to the mathematical statement is needed. It does not close a narrower question the todobox did not originally ask but that a careful reader might: the quoted sentence is Papadopoulos's English translation of Tissot's French, not a transcription of Tissot's original wording in French, and this chapter has not itself opened a scanned copy of pages 1–21 of the 1881 *Mémoire* (Internet Archive item `mmoiresurlarepr00tissgoog`, cited in this section's opening footnote) to check the French original word for word. Given that Papadopoulos's account is a specialist's direct reading with specific page citations, corroborated by his statement that the 1881 text's first chapter proves rather than revises the 1859 note's claim, this is treated here as a substantial upgrade over the previous tertiary sourcing, not as a fully closed bibliographic question.

## 3.2 Setup: the local linearization of a projection

Let $f$ be a map projection, i.e. a smooth map from a region of the sphere (parametrized, as in chapter 2, by $(\varphi, \lambda)$ with metric $ds^2 = R^2\, d\varphi^2 + R^2 \cos^2\varphi\, d\lambda^2$) to the plane, with plane coordinates $(x, y) = f(\varphi, \lambda)$. Fix a point $P_0 = (\varphi_0, \lambda_0)$. Because $f$ is smooth, it is well approximated near $P_0$ by its differential $df_{P_0}$, a linear map from the tangent plane of the sphere at $P_0$ to $\mathbb{R}^2$. Tissot's construction is simply the geometric content of this linear approximation: a linear map takes circles to ellipses, and the shape of that ellipse, at each point, is the entire local distortion signature of the projection at that point.

Concretely, define the local orthonormal frame on the sphere at $P_0$,

$$\mathbf{e}_\varphi = \frac{1}{R}\frac{\partial \mathbf{x}}{\partial \varphi}, \qquad \mathbf{e}_\lambda = \frac{1}{R\cos\varphi}\frac{\partial \mathbf{x}}{\partial \lambda},$$

so that a unit tangent vector at angle $\theta$ from the meridian is $\mathbf{u}(\theta) = \cos\theta\,\mathbf{e}_\varphi + \sin\theta\,\mathbf{e}_\lambda$. Under $df_{P_0}$, this unit vector is sent to

$$df_{P_0}(\mathbf{u}(\theta)) = \cos\theta \cdot \frac{1}{R}\frac{\partial(x,y)}{\partial\varphi} + \sin\theta \cdot \frac{1}{R\cos\varphi}\frac{\partial(x,y)}{\partial\lambda},$$

a vector in the plane whose length, as $\theta$ ranges over $[0, 2\pi)$, traces out an ellipse: the *Tissot indicatrix* at $P_0$.

## 3.3 Principal directions and principal scales

**Theorem 3.1** (Tissot, 1859/1881)**.** *At every point $P_0$ where $df_{P_0}$ is non-singular, there exist two orthogonal directions on the sphere — the* principal directions *— whose images under $df_{P_0}$ are also orthogonal in the plane. Along these two directions, and only these two (generically), the local length-scaling factor $ds'/ds$ attains its maximum value $a$ and minimum value $b$ respectively; the indicatrix is precisely the ellipse with semi-axes $a, b$ aligned with the images of the principal directions.*

*Proof sketch.* This is exactly the statement that the symmetric positive-definite $2 \times 2$ matrix $M = (df_{P_0})^\top df_{P_0}$, expressed in the orthonormal frame $\{\mathbf{e}_\varphi, \mathbf{e}_\lambda\}$, is diagonalizable by an orthogonal change of basis — the spectral theorem for symmetric matrices, applied to the pulled-back Euclidean metric on the tangent plane. The eigenvalues of $M$ are $a^2$ and $b^2$; the eigenvectors give the principal directions. Because $M$ is exactly the local representation of the first fundamental form of the target metric pulled back through $f$, this is the same mathematical fact used, at the level of the metric tensor itself rather than a single point's linearization, in the definition of $E, F, G$ in chapter 2. □

In the classical cartographic literature, and throughout the computational chapters of *No Map Is Best* and the DGGS monograph, the principal scale factors $a, b$ are usually computed directly, without explicitly diagonalizing $M$, via the formulas (for a projection $x(\varphi,\lambda), y(\varphi,\lambda)$)

$$h = \frac{1}{R}\sqrt{x_\varphi^2 + y_\varphi^2}, \qquad k = \frac{1}{R\cos\varphi}\sqrt{x_\lambda^2 + y_\lambda^2}, \qquad \sin\varepsilon = \frac{x_\varphi y_\lambda - y_\varphi x_\lambda}{R^2\cos\varphi\, hk},$$

where $h, k$ are the scale factors along the meridian and parallel respectively and $\varepsilon$ is the angle between their images (which need not be 90° even though meridians and parallels are orthogonal on the sphere itself — this non-orthogonality is itself a form of distortion). The principal scales $a \geq b$ then follow from

$$a + b = \sqrt{h^2 + k^2 + 2hk\sin\varepsilon}, \qquad a - b = \sqrt{h^2 + k^2 - 2hk\sin\varepsilon}. \tag{3.1}$$

## 3.4 The three canonical distortion measures, read off the indicatrix

**Key Result**

*Every scalar distortion measure used elsewhere in this series is a specific function of the single pair $(a, b)$:*

- ***Areal scale.*** *The area of the image ellipse relative to the area of the source unit circle is $s = ab$. A projection is* equal-area *(equivalent) at $P_0$ if and only if $ab = 1$.*
- ***Maximum angular deformation.*** *The greatest possible difference between an angle on the sphere and the corresponding angle on the map, attained by directions bisecting the principal axes, is*

$$\omega = 2 \arcsin\left(\frac{a-b}{a+b}\right).$$

*A projection is* conformal *at $P_0$ if and only if $a = b$ (the indicatrix is a circle, $\omega = 0$).*

- ***Linear scale in any direction*** *$\theta$ (measured from the major axis of the indicatrix) is $\sqrt{a^2 \cos^2\theta + b^2 \sin^2\theta}$, interpolating between the two principal extremes.*

Corollary 2.2 showed that $a = b = 1$ everywhere is impossible on any open region. Tissot's construction is what makes that abstract impossibility *quantitative*: it is not merely true that some projection property must fail, but the indicatrix tells us, at every single point and for every specific projection formula, exactly which properties fail and by how much. This is the sense in which this chapter is the working bridge between chapter 2's existence theorem and the computational chapters of the companion volumes: the Tissot ellipse eccentricity computed by finite differences in DGGS §§(Tissot indicatrix chapter) and the areal-scale figures computed projection-by-projection in *No Map Is Best* are both direct numerical instances of $(a, b)$ as defined here.

## 3.5 The indicatrix as a manifestation of non-zero curvature

It is worth making the connection to chapter 2 completely explicit, since the two chapters are sometimes presented, in expository treatments, as independent topics rather than as two views of one fact. Consider a conformal projection ($a = b \equiv \mu(\varphi, \lambda)$ everywhere, the scale factor $\mu$ possibly varying with position). Chapter 5 carries out, in full, the substitution into the orthogonal-coordinates curvature formula eq. (2.4) that this section only sketched in an earlier draft; the resulting identity is

$$\Delta_g(\ln \mu) = K,$$

where $\Delta_g$ is the *Laplace–Beltrami operator of the sphere's own metric* $ds^2 = R^2 d\varphi^2 + R^2 \cos^2\varphi \, d\lambda^2$ — not the flat coordinate Laplacian — and $K \equiv 1/R^2$ is the sphere's constant curvature. Since $K \neq 0$ (chapter 2), $\ln \mu$ can never be $\Delta_g$-harmonic on any open region: $\mu$ can never be constant there. This differential identity — that the log-scale factor of any conformal map of the sphere is forced to satisfy a genuinely *linear* Poisson equation, $\Delta_g(\ln \mu) = K$, with the curved surface's own Laplace–Beltrami operator on the left and a constant source on the right — is the seed of the Milnor-theorem chapter (chapter 5), where it is turned into both a global flux bound and a pointwise maximum-principle bound.

**Epistemic Notice**

An earlier draft of this section quoted, from memory of the standard isothermal-coordinates formula $K = -\lambda^{-2}\Delta_0(\ln\lambda)$ (flat $\Delta_0$, $\lambda$ the conformal factor *of the source metric relative to a flat reference chart*), a superficially similar but materially different identity, $K = -\mu^{-2}\Delta(\ln\mu)$ with $\Delta$ read as a flat coordinate Laplacian acting directly on $\mu$, the *projection's* scale factor, rather than on $\lambda$, the sphere's own isothermal conformal factor. Chapter 5 carries out the substitution into eq. (2.4) explicitly rather than quoting the standard formula by analogy, and the two are not the same computation: $\mu$ compares the image metric to the sphere metric (both generally curved, in the projection's own coordinates), not the sphere metric to a flat reference chart, and the correct identity for $\mu$ is linear in $\ln\mu$, with the sphere's own Laplace–Beltrami operator, not the semilinear flat-Laplacian form originally stated here. This is recorded as an explicit correction, not silently amended, precisely because theorem 5.2 in chapter 5 was built on the original (incorrect) form; the derivation given in chapter 5 is self-contained and verified independently of this section's original claim, including against the textbook tangent-stereographic scale formula $\mu(\theta) = \sec^2(\theta/2)$ as a concrete check.

# Part II

# The Impossibility Theorems Proper

# Chapter 4

# No Map Is Both Conformal and Equal-Area

## 4.1 Statement, and why it needs a chapter of its own

The result proved in this chapter is, formally, a two-line corollary of material already in hand: theorem 3.1 showed that every projection's local distortion at a point is completely described by the principal scales $(a, b)$; conformality was defined there as $a = b$, equal-area as $ab = 1$. Combining the two forces $a = a \cdot a = ab = 1$, hence $a = b = 1$: an isometry at that point, which corollary 2.2 already rules out except on a measure-zero (indeed empty, for an open region) set. So why does this deserve to be separated out as item (i) of the introduction's list, with its own chapter, rather than folded into a remark at the end of chapter 3?

Three reasons, each of which the rest of the chapter develops. First, the two-line argument proves the claim *pointwise*, but the cartographic literature, and the companion volumes in this series, use "conformal" and "equal-area" as properties of a projection formula holding on an *entire region* at once, and the passage from a pointwise impossibility to a global one, while immediate here, is exactly the place where the next chapter's argument (Milnor) is not immediate, so it is worth being careful about the logical structure in the easy case first. Second, the pointwise statement admits a sharp converse that is cartographically important in its own right and is not simply the negation of the impossibility: at every point, for every value of the areal-scale target, there is a definite, computable amount of angular distortion that must be accepted, and conversely; this quantitative trade-off, not just its qualitative existence, is what section 4.2 derives. Third, and most importantly for the volume's stated aim of unifying material already scattered across the companion series, the theorem organizes essentially every named projection appearing in *No Map Is Best* and the DGGS monograph into exactly one of three mutually exclusive families, and section 4.4 makes that organization explicit for the first time, seeding the full taxonomy planned for chapter 26.

**Theorem 4.1** (Impossibility of simultaneous conformality and equivalence)**.** *Let $f$ be a smooth projection defined on a non-empty open region $U$ of the sphere, with principal scales $a(P) \geq b(P) > 0$ at each $P \in U$ as in theorem 3.1. If $f$ is conformal on $U$ (i.e. $a(P) = b(P)$ for every $P \in U$) and equal-area on $U$ (i.e. $a(P)b(P) = 1$ for every $P \in U$), then $a(P) = b(P) = 1$ for every $P \in U$, i.e. $f$ is a local isometry on $U$. By corollary 2.2 no such $f$ exists. Hence no projection of a non-empty open region of the sphere is simultaneously conformal and equal-area.*

*Proof.* Immediate substitution, as above: conformality gives $a = b$ at each point; equivalence gives $ab = 1$ at each point; together, $a^2 = 1$, and since $a > 0$ by non-singularity of $df_P$ (theorem 3.1), $a = 1$, hence $b = 1$. A projection with $a \equiv b \equiv 1$ on $U$ preserves the length

of every tangent vector at every point of $U$, hence (by integrating along paths) the length of every rectifiable curve in $U$: this is precisely the definition of a local isometry from chapter 2. Corollary 2.2 excludes this for any non-empty open $U$. □

This is worth contrasting immediately with what the theorem does *not* say, because the two failure modes it rules out are easy to conflate with two related but strictly weaker properties that *are* individually achievable everywhere.

*Remark* 4.2 (Global conformality and global equivalence are each separately possible). Theorem 4.1 forbids the conjunction, not either conjunct on its own. A projection can be conformal on the whole sphere minus a point (stereographic projection, $a = b = \sec^2(\text{colatitude}/2)$ pointwise, never equal to a constant but always with $a = b$); a projection can be equal-area on the whole sphere (the Lambert azimuthal equal-area projection, or the sinusoidal, or Snyder's Eckert-family constructions surveyed projection-by-projection in *No Map Is Best*). What is impossible is a single formula lying in *both* families on any open set, however small.

## 4.2 The trade-off curve at a point

Theorem 4.1 is a statement that a certain pair of constraints is jointly infeasible. It is natural to ask what *is* feasible: given that a cartographer at a specific point $P_0$ is willing to accept some fixed amount $s > 0$ of areal distortion (meaning $ab = s$, not necessarily 1), how little angular distortion can be bought, and vice versa? This is answered completely by eq. (3.1) together with the definitions in chapter 3, and the answer is a genuine two-parameter family, not a single forced value, which is what makes cartographic design a real design problem rather than a solved one.

**Proposition 4.3** (Feasible region in $(s, \omega)$-space)**.** *Fix a point $P_0$. As the pair $(a, b)$ ranges over all values achievable by* some *smooth projection at $P_0$ (i.e. over all $a \geq b > 0$, since any positive-definite symmetric $2 \times 2$ matrix with any eigenvalues is realized by some local diffeomorphism), the areal scale $s = ab$ and the maximum angular deformation $\omega = 2\arcsin\left(\frac{a-b}{a+b}\right)$ satisfy exactly the constraint*

$$0 \leq \omega < \pi, \qquad s > 0,$$

*with no further relation between them: for every pair $(s, \omega)$ in this open range there exists $a \geq b > 0$, unique up to the trivial symmetry $s$ fixed, realizing it. In particular the level set $\{\omega = 0\}$ (conformal) meets the level set $\{s = 1\}$ (equal-area) only at the single point $a = b = 1$, consistently with theorem 4.1, and away from that single point, decreasing $\omega$ toward 0 while holding $s \neq 1$ fixed is always achievable, i.e. one can get* arbitrarily close *to conformal while remaining exactly equal-area everywhere except at isolated points or curves, which is exactly the compromise strategy used by, e.g., the Winkel Tripel and other "minimum-error" compromise projections surveyed in* No Map Is Best.

*Proof.* Given target $(s, \omega)$ with $s > 0$ and $0 \leq \omega < \pi$, set $r = \sin(\omega/2) = \frac{a-b}{a+b} \in [0, 1)$ and solve the two equations $ab = s$, $(a - b)/(a + b) = r$ simultaneously: from the second, $a = b\,\frac{1+r}{1-r}$, substitute into the first to get $b^2\frac{1+r}{1-r} = s$, i.e. $b = \sqrt{s\,\frac{1-r}{1+r}}$, and then $a = \sqrt{s\,\frac{1+r}{1-r}}$; both are well defined and positive since $r < 1$ and $s > 0$, giving the unique pair $(a, b)$ with $a \geq b$. Conversely $\omega = \pi$ would require $r = 1$, i.e. $b = 0$, excluded since $df_{P_0}$ is required non-singular. □

**Key Result**

*Proposition 4.3 is the precise sense in which theorem 4.1 is a statement about a single forbidden point in an otherwise open, two-dimensional space of trade-offs, not a statement that distortion is large or that design freedom is small. Every numerical areal-scale and angular-deformation figure computed projection-by-projection in* No Map Is Best *is a sample point* $(s, \omega)$ *from exactly this feasible region, and the region's only excluded point is* $(1, 0)$.

## 4.3 From a point to a region: why the passage is not vacuous

Theorem 4.1 as stated already applies to an entire open region $U$, not just a point, because the hypothesis was that $a(P) = b(P)$ and $a(P)b(P) = 1$ hold *at every* $P \in U$ simultaneously, and the pointwise argument then applies at each $P$ separately; no integration or global argument was needed. It is worth flagging explicitly that this is a special feature of *this* theorem and will not recur in chapter 5, where the quantity bounded is a genuinely global (integrated) one and the pointwise inequality it rests on is strict except possibly at isolated points, so that the interesting content of the theorem is exactly in how the pointwise slack accumulates over the region. Making this contrast explicit here is intended to prevent the reader from expecting the same easy point-to-region passage to work again in the next chapter.

## 4.4 Taxonomic consequences: three families, no fourth

**Definition 4.4** (Projection families by local behaviour)**.** A projection $f$ on a region $U$ is:

- *conformal* on $U$ if $\omega(P) \equiv 0$ (equivalently $a(P) = b(P)$) for all $P \in U$;
- *equal-area* (equivalent, or *authalic*) on $U$ if $s(P) \equiv 1$ (equivalently $a(P)b(P) = 1$) for all $P \in U$;
- *aphylactic* (or *compromise*) on $U$ if neither of the above holds identically on $U$, so that $0 < \omega(P) < \pi$ and $s(P) \neq 1$ hold at least somewhere (possibly with isolated exceptional points, such as the standard parallels of a conic compromise projection, where one or the other property holds only locally).

**Corollary 4.5** (Exhaustiveness and exclusivity of the three families)**.** *Every projection on a non-empty open* $U$ *belongs to at least one of the three families above (exhaustiveness, immediate from the trichotomy on* $(\omega, s)$ *at each point extended to the region by the evident case split). No projection on a non-empty open* $U$ *belongs to both the conformal and equal-area families (exclusivity, this is exactly theorem 4.1). A projection can belong to the aphylactic family and to neither of the other two on all of* $U$ *(e.g. the Miller cylindrical, or any of the compromise constructions in* No Map Is Best*), or it can fail to be conformal or equal-area on* $U$ *as a whole while still achieving one property exactly along a lower-dimensional subset of* $U$ *(a curve of true scale, a pair of standard parallels) — this intermediate case is the one that motivates recording, rather than discarding, the specific curve or point set on which a named projection achieves exact conformality or equivalence, a bookkeeping practice already followed systematically for each projection profiled in* No Map Is Best.

**Epistemic Notice**

The three-family taxonomy of corollary 4.5 is standard in the cartographic literature (it corresponds to what Snyder 1987 and the USGS professional paper series call conformal, equal-area, and "compromise" or "aphylactic" projections) and is not a new classification; what is new to this chapter is deriving the trichotomy directly from theorems 3.1 and 4.1 rather than presenting it, as the applied cartographic literature typically does, as an empirical grouping of known named projections. The full closing taxonomy planned for chapter 26 extends corollary 4.5 by a second, independent axis (global versus locally/centrally defined, i.e. azimuthal-type constructions such as the retroazimuthal projections of the companion volume) and argues that the two axes together, not the areal/angular axis alone, are what is needed to place every projection treated anywhere in this series into a single exhaustive grid; that two-axis argument is not made here.

Corollary 4.5's original todobox flagged the exceptional-point-set case as asserted without proof. Proposition 4.6 immediately below supplies the general argument (regular level set theorem), and the disclaimerbox following its proof states the remaining caveat explicitly (regularity is not verified case-by-case for every named projection). Section 4.5 below closes that caveat for the two conic families this series actually uses, and in the process finds and fixes a genuine gap in proposition 4.6 itself (proposition 4.7). This box records that the chain is complete within this chapter; no further action is needed on this specific item.

**Proposition 4.6** (Exceptional sets are curves, away from critical points)**.** *Let $f$ be a smooth projection on $U$ with $\omega, s : U \to \mathbb{R}$ as in proposition 4.3, and suppose $f$ is not conformal (respectively not equal-area) on all of $U$. If $0$ is a regular value of $\omega$ (respectively $1$ is a regular value of $s$) — i.e. $\nabla\omega \neq 0$ (resp. $\nabla s \neq 0$) at every point where $\omega = 0$ (resp. $s = 1$) — then the exceptional set $\{\omega = 0\}$ (resp. $\{s = 1\}$) is a smooth $1$-dimensional embedded submanifold of $U$: locally, a curve.*

*Proof.* This is the regular level set theorem applied directly: $\omega : U \to \mathbb{R}$ is a smooth map from a 2-dimensional manifold to $\mathbb{R}$, so at any point $P_0$ with $\omega(P_0) = 0$ and $\nabla\omega(P_0) \neq 0$, the implicit function theorem gives local coordinates in which $\omega$ is one of the two coordinate functions, so $\{\omega = 0\}$ is locally the graph of a smooth function of the other coordinate — a 1-dimensional curve through $P_0$. Since this holds at every point of $\{\omega = 0\}$ by the regular-value hypothesis, the whole set is a smooth embedded 1-manifold (a disjoint union of such local curve pieces). The argument for $\{s = 1\}$ is identical with $s - 1$ in place of $\omega$. □

**Epistemic Notice**

This closes the gap identified in the previous todobox, but only conditionally, and the condition is genuinely doing work, not a formality: the proposition assumes 0 (resp. 1) is a *regular* value, and this has not been shown to hold for every named projection profiled in *No Map Is Best*, only assumed plausible by analogy with the standard-parallel examples where it visibly does hold (a conic projection's standard parallels are, generically, where $\partial\omega/\partial\varphi \neq 0$, by the same computation that shows $\omega$ increases away from them). At a *critical* point of $\omega$ on $\{\omega = 0\}$ — a point where the projection is conformal to first order in a stronger sense, an inflection of the distortion field — the level set can fail to be a manifold (it can have a self-crossing, an isolated point, or worse), and proposition 4.6 says nothing about that case. No claim is made here that regularity holds universally across the series' named projections; verifying it case by case (or finding a counterexample) is left for a future revision, distinct from the general argument now on record.

## 4.5 Closing the conic case, and a correction proposition 4.6 itself needs

The todobox above singled out the conic standard-parallel case as the plausible instance left unverified, and chapter 16 later verified the transverse-cylindrical analogue for Gauss–Krüger while explicitly noting that the conic case itself was still open. This section closes it, for the two named conic families this series actually uses — the Lambert Conformal Conic (LCC, a member of the conformal family, whose exceptional set is $\{s = 1\}$) and the Albers Equal-Area Conic (a member of the equal-area family, whose exceptional set is $\{\omega = 0\}$) — and in the course of doing so finds that proposition 4.6 as stated has a gap of its own, not only an unverified hypothesis: at exactly the points it is meant to describe, $\omega$ need not be differentiable at all, so "0 is a regular value of $\omega$" is not merely a strong hypothesis but, for the equal-area family, generically a false one. The fix is a small one, using material already in hand.

**Key Result**

*[Why $\omega$ can fail to be smooth on its own zero set] Write $\omega = 2\arcsin\big(|h-k|/(h+k)\big)$ for two smooth positive functions $h, k$ on $U$ (the meridian and parallel scale factors of theorem 3.1, with $a = \max(h,k)$, $b = \min(h,k)$, matching proposition 4.3's formula since* arcsin *is applied to a manifestly non-negative, order-independent quantity). If $g := h - k$ has a simple (transversal) zero at $P_0 \in U$, i.e. $g(P_0) = 0$ and $\nabla g(P_0) \neq 0$, then $\omega = 2\arcsin(|g|/(h+k))$ is, near $P_0$, the composition of a smooth, strictly increasing odd function of $|g|/(h+k)$ with $|g|$ itself, and $|g|$ has unequal one-sided directional derivatives along $\nabla g(P_0)/|\nabla g(P_0)|$ (namely $+|\nabla g(P_0)|$ and $-|\nabla g(P_0)|$) at $P_0$. Consequently $\omega$ is continuous but* not *differentiable at $P_0$: it has a genuine corner there, with slope $\pm 2|\nabla g(P_0)|/(h(P_0)+k(P_0))$ on the two sides. This holds regardless of which specific projection $h, k$ come from; it is a structural fact about the definition of $\omega$ itself, not an artifact of the conic examples below. The one case it does* not *apply to is $\nabla g(P_0) = 0$ — $h$ and $k$ tangent rather than transversal at $P_0$ — where $\omega$ can be smooth (or not) depending on the higher-order behaviour of $g$, and section 4.5 below is exactly such a case.*

**Proposition 4.7** (Corrected sufficient condition for proposition 4.6)**.** *Let $f$ be conformal on $U$ (so $h \equiv k$ and $s = k^2$; the exceptional set of interest is $\{s = 1\}$), or equal-area on $U$ (so $hk \equiv 1$; the exceptional set of interest is $\{\omega = 0\} = \{h = k\}$). In the conformal case, if 1 is a regular value of the smooth function $s = k^2$, proposition 4.6's conclusion holds verbatim: $\{s = 1\}$ is a smooth embedded curve. In the equal-area case, replace "0 is a regular value of $\omega$" by the strictly weaker and always well-posed condition "0 is a regular value of the smooth function $g := h - k$"; under this condition $\{g = 0\} = \{\omega = 0\}$ is a smooth embedded curve by the ordinary regular level set theorem applied to $g$ (not to $\omega$, which need not even be differentiable there by section 4.5), and this is the correct reading of what proposition 4.6 needed all along for the equal-area family.*

*Proof.* Conformal case: identical to proposition 4.6's own proof with $s$ in place of $\omega$, since $s$ is smooth everywhere on $U$ (no absolute value is involved: $s = k^2$ directly, with no case split on the ordering of $h, k$). Equal-area case: $g = h - k$ is smooth on all of $U$ by hypothesis (both $h, k$ are smooth), so the regular level set theorem applies to $g$ directly wherever $\nabla g \neq 0$, giving a smooth embedded curve $\{g = 0\}$; and $\{g = 0\} = \{h = k\} = \{\omega = 0\}$ as sets, by proposition 4.3's formula for $\omega$, regardless of whether $\omega$ itself is differentiable along that set. □

We now verify the hypothesis of proposition 4.7 for both named conic families, using the standard spherical Snyder formulas for the two constructions (secant form, two distinct standard parallels $\varphi_1 \neq \varphi_2$, the form actually deployed — e.g. the 29.5°/45.5° pair used for

the USGS Albers map of the contiguous United States) rather than the degenerate tangent (single standard parallel) form.

**Proposition 4.8** (Regularity at LCC standard parallels)**.** *For the spherical Lambert Conformal Conic with two distinct standard parallels $\varphi_1 \neq \varphi_2$ in $(-\pi/2, \pi/2)$, cone constant $n = \dfrac{\ln(\cos\varphi_1/\cos\varphi_2)}{\ln\big(\tan(\pi/4+\varphi_2/2)/\tan(\pi/4+\varphi_1/2)\big)}$, and parallel scale factor $k(\varphi) = \dfrac{n\rho(\varphi)}{R\cos\varphi}$ (equal to the meridian scale factor $h(\varphi)$ everywhere, since the construction is conformal by design), the exceptional set $\{s = 1\} = \{k = 1\}$ consists exactly of the two parallels $\varphi = \varphi_1, \varphi_2$, and at each of them*

$$\left.\frac{dk}{d\varphi}\right|_{\varphi=\varphi_i} = \frac{\sin\varphi_i - n}{\cos\varphi_i},$$

*which is nonzero for both $i = 1, 2$ whenever $\varphi_1 \neq \varphi_2$.*

*Proof.* Write $t(\varphi) = \tan(\pi/4 + \varphi/2)$, so $\rho = RFt^{-n}$ for a constant $F$ chosen so $k(\varphi_1) = 1$, and $t'(\varphi)/t(\varphi) = \sec\varphi$ (a direct computation: $t' = \frac{1}{2}\sec^2(\pi/4 + \varphi/2)$, and $2\sin(\pi/4 + \varphi/2)\cos(\pi/4 + \varphi/2) = \sin(\pi/2 + \varphi) = \cos\varphi$, so $t'/t = \frac{1}{2}\sec^2(\cdot)/\tan(\cdot) = \frac{1}{2\sin(\cdot)\cos(\cdot)} = 1/\cos\varphi$). Then $\ln k = \text{const} - n\ln t - \ln\cos\varphi$, so $\dfrac{d\ln k}{d\varphi} = -n\sec\varphi + \tan\varphi$; at a point where $k = 1$ this equals $dk/d\varphi$ directly, giving the stated formula. That $k(\varphi_1) = k(\varphi_2) = 1$ exactly, by construction, is the defining property of $F$ and $n$ (section 4.5 below re-derives it independently). For the nonvanishing claim: apply the Cauchy Mean Value Theorem to $f(\varphi) = \ln\cos\varphi$ and $g(\varphi) = \ln t(\varphi)$ on $[\varphi_1, \varphi_2]$ (or $[\varphi_2, \varphi_1]$); $g$ is strictly monotonic since $g' = t'/t = \sec\varphi > 0$ throughout $(-\pi/2, \pi/2)$, so the theorem applies validly and gives $\xi$ strictly between $\varphi_1$ and $\varphi_2$ with

$$n = \frac{f(\varphi_1) - f(\varphi_2)}{g(\varphi_2) - g(\varphi_1)} = -\frac{f'(\xi)}{g'(\xi)} = -\frac{-\tan\xi}{\sec\xi} = \sin\xi.$$

Since sin is strictly increasing on $(-\pi/2, \pi/2)$ and $\xi$ lies strictly between $\varphi_1, \varphi_2$, $n = \sin\xi$ lies strictly between $\sin\varphi_1$ and $\sin\varphi_2$, hence $n \neq \sin\varphi_1$ and $n \neq \sin\varphi_2$. The displayed derivative is therefore nonzero at both standard parallels. □

**Proposition 4.9** (Regularity at Albers standard parallels)**.** *For the spherical Albers Equal-Area Conic with two distinct standard parallels $\varphi_1 \neq \varphi_2$, $n = \frac{1}{2}(\sin\varphi_1 + \sin\varphi_2)$, $C = \cos^2\varphi_1 + 2n\sin\varphi_1$, $\rho(\varphi) = R\sqrt{C - 2n\sin\varphi}/n$, and $k(\varphi) = n\rho(\varphi)/(R\cos\varphi)$ (with $h \equiv 1/k$ everywhere, since the construction is equal-area by design), the exceptional set $\{\omega = 0\} = \{h = k\} = \{k = 1\}$ again consists exactly of $\varphi = \varphi_1, \varphi_2$, and*

$$\left.\frac{dk}{d\varphi}\right|_{\varphi=\varphi_i} = \frac{\sin\varphi_i - n}{\cos\varphi_i} = \pm\frac{\sin\varphi_2 - \sin\varphi_1}{2\cos\varphi_i},$$

*nonzero for both $i$ whenever $\varphi_1 \neq \varphi_2$; consequently $g = h - k = 1/k - k$ satisfies $dg/d\varphi|_{\varphi_i} = -(1 + 1/k^2)\,dk/d\varphi|_{\varphi_i} = -2\,dk/d\varphi|_{\varphi_i} \neq 0$ as well, so proposition 4.7's hypothesis holds.*

*Proof.* At $\varphi = \varphi_1$: $C - 2n\sin\varphi_1 = \cos^2\varphi_1 + 2n\sin\varphi_1 - 2n\sin\varphi_1 = \cos^2\varphi_1$, so $\rho(\varphi_1) = R\cos\varphi_1/n$ and $k(\varphi_1) = n\cdot(R\cos\varphi_1/n)/(R\cos\varphi_1) = 1$; the same computation at $\varphi_2$ uses $C - 2n\sin\varphi_2 = \cos^2\varphi_1 + 2n(\sin\varphi_1 - \sin\varphi_2) = \cos^2\varphi_1 - (\sin\varphi_2 - \sin\varphi_1)^2$, which a short direct expansion (using $2n = \sin\varphi_1 + \sin\varphi_2$) shows equals $\cos^2\varphi_2$, giving $k(\varphi_2) = 1$ as well. Differentiating $\ln k = \frac{1}{2}\ln(C - 2n\sin\varphi) - \ln\cos\varphi$ gives $\dfrac{d\ln k}{d\varphi} = \dfrac{-n\cos\varphi}{C - 2n\sin\varphi} + \tan\varphi$; at $\varphi_1$, where $C - 2n\sin\varphi_1 = \cos^2\varphi_1$, this is $-n/\cos\varphi_1 + \tan\varphi_1 = (\sin\varphi_1 - n)/\cos\varphi_1$, and since $k = 1$ there this equals $dk/d\varphi|_{\varphi_1}$ directly; the identical computation at $\varphi_2$ (using

$C - 2n\sin\varphi_2 = \cos^2\varphi_2$) gives the same formula with $\varphi_2$ in place of $\varphi_1$. Since $n$ is the exact arithmetic mean of $\sin\varphi_1, \sin\varphi_2$ by definition, $n - \sin\varphi_1 = \frac{1}{2}(\sin\varphi_2 - \sin\varphi_1)$ and $n - \sin\varphi_2 = \frac{1}{2}(\sin\varphi_1 - \sin\varphi_2)$, both nonzero exactly when $\varphi_1 \neq \varphi_2$, which is assumed. □

**Epistemic Notice**

The Snyder formulas quoted above are the standard spherical-model textbook forms for these two named projections (Snyder, *Map Projections: A Working Manual*, USGS Professional Paper 1395, 1987, already cited elsewhere in this series' taxonomy discussion, corollary 4.5's own disclaimer-box), not a new derivation; what is new here is the differentiation and the mean-value argument built on top of them. All four claims above — $k(\varphi_1) = k(\varphi_2) = 1$ exactly, the closed-form derivative, its nonvanishing, and (for Albers specifically) $\omega$'s failure to be differentiable at the very points proposition 4.6 concerns — were checked independently in floating point (40 decimal digits, `mpmath`) against four different $(\varphi_1, \varphi_2)$ pairs spanning both hemispheres and both narrow and wide parallel separations, by `scripts/conic_standard_parallel_regularity.py`, and the degenerate tangent limit $\varphi_2 \to \varphi_1$ was checked separately and confirmed to drive $dk/d\varphi \to 0$ at three further test latitudes, matching section 4.5 below.

**Key Result**

*[The tangent (single standard parallel) case is genuinely different, not merely a limit worth skipping] In the degenerate limit $\varphi_2 \to \varphi_1 = \varphi_0$ (a cone, or in the further limit $\varphi_0 \to 0$ a cylinder, tangent along a single parallel rather than secant along two), $n \to \sin\varphi_0$ exactly, so propositions 4.8 and 4.9's closed form gives $dk/d\varphi|_{\varphi_0} = (\sin\varphi_0 - \sin\varphi_0)/\cos\varphi_0 = 0$: the tangent standard parallel is a* critical*, not regular, point of $k$ (and hence of $s = k^2$ for LCC, and of $g = h - k$ for Albers), confirmed independently above to second order (the numerically computed $d^2k/d\varphi^2$ at the tangent parallel is 1, matching the classical Mercator fact $k(\varphi) = \sec\varphi$, $k''(0) = 1$, itself the cylindrical $\varphi_0 = 0$ special case of exactly this same phenomenon). Proposition 4.7 accordingly does* not *apply to the tangent case, for either family — but the exceptional set is nonetheless a smooth curve there too (indeed a single parallel circle, the cleanest possible curve), by an entirely different argument that owes nothing to the regular level set theorem: $k$ (or $g$) depends on $\varphi$ alone, not on longitude, so its full gradient in $(\varphi, \lambda)$ coordinates vanishes identically along the whole parallel $\varphi = \varphi_0$ whenever it vanishes at one point of it, and a level set of a function depending on a single coordinate is trivially that coordinate's fiber, smooth regardless of whether the function's derivative in that one variable happens to vanish there too. This is worth recording explicitly because it shows the regular-value hypothesis in propositions 4.6 and 4.7 is sufficient but never necessary: the tangent conic and cylindrical cases are counterexamples to necessity that do not touch the propositions' correctness, and the secant conic cases of propositions 4.8 and 4.9 are where the hypothesis is both needed and, now, verified rather than assumed.*

**Epistemic Notice**

This closes, for the two conic families this series actually names and uses (LCC and Albers, each in the secant, two-standard-parallel form that is the one actually deployed in named historical and administrative maps), the specific gap chapter 16 left open after closing the transverse-cylindrical instance. It does not close the general claim for every named projection profiled in *No Map Is Best*; the corrected statement, proposition 4.7, is what any future case-by-case check should verify against, and section 4.5's general fact about $\omega$'s non-smoothness on its own zero set applies to every equal-area-family exceptional-set check in that future work, not only to Albers, so it is recorded here once rather than separately at each such site.

# Chapter 5

# From Pointwise to Global: A Gauss–Bonnet Bound, and Milnor's Theorem

**Epistemic Notice**

This chapter is the most epistemically hazardous in the volume so far, and it is worth saying so plainly before it begins. Section 5.2 derives a genuine, self-contained inequality (now, after the correction recorded in that section, an honest theorem proved here in full from eq. (2.4) of chapter 2) bounding a global distortion functional in terms of total curvature. Section 5.6, by contrast, describes John Milnor's 1969 result (*A Problem in Cartography*, *American Mathematical Monthly* 76(10), pp. 1101–1112) at the level of its problem statement and the qualitative shape of its conclusion only. The precise functional Milnor minimizes, the exact value of the sharp bound, and the claimed uniqueness and identity of the extremal projection are stated here from secondary recollection of the cartographic literature that cites Milnor's result (chiefly Snyder's discussions of "minimum-error" projections), *not* from a direct reading of the primary 1969 paper, and every such claim is explicitly flagged and boxed as unverified. This is a stronger caveat than the source-provenance notes in chapters 2 and 3: those chapters' central theorems were derived in full here regardless of citation quality; this chapter's title theorem is not.

## 5.1 What a "global" distortion bound should even mean

Chapter 3 gave the pointwise distortion pair $(a(P), b(P))$ at each point, and chapter 4 showed that the two pointwise targets $a = b$ and $ab = 1$ cannot both be hit at any single point of an open region. But a cartographer choosing a projection for, say, a hemisphere does not evaluate it at one point; the practical question is how the distortion *accumulates* or *averages* over the whole mapped region, and in particular whether making distortion small at the centre of the map (as every named projection in *No Map Is Best* does, by construction, at its central point or standard parallel) forces it to be correspondingly large somewhere else. This is a different, and harder, question than the pointwise one, because in principle a projection could in principle trade a small distortion at the centre for an unboundedly large one only on a shrinking neighbourhood of the boundary, with the two effects failing to show up in any naive average. Ruling this out — showing that some fixed, non-negligible amount of distortion is forced *in total*, not just somewhere — is the content both of the bound derived below and of Milnor's actual theorem.

## 5.2 A genuine bound: total curvature controls total conformal defect

Restrict attention, for this section, to *conformal* projections, $a \equiv b \equiv \mu(\varphi, \lambda)$. Chapter 3 quoted, from memory of the standard isothermal-coordinates formula, an identity relating $\mu$ to the sphere's curvature $K \equiv 1/R^2$, with a promise that the full derivation, from the orthogonal-coordinates curvature formula eq. (2.4) of chapter 2, would be carried out here. Carrying it out reveals that the quoted identity was stated incorrectly — in both sign and functional form — and this section corrects it before building anything further on it.

**Epistemic Notice**

[Correction to a previously quoted identity] Chapter 3 stated $K = -\mu^{-2}\Delta(\ln \mu)$ with $\Delta$ read as the flat coordinate Laplacian acting on the projection's own scale factor $\mu$, by direct (but, it turns out, mistaken) analogy with the standard isothermal-coordinates formula for a conformal factor *relative to a flat reference chart*. That analogy does not transfer: $\mu$ is the ratio of the image metric to the *sphere's own* metric, and the sphere's own metric is not flat, so the flat-Laplacian formula does not apply to $\mu$ directly. The derivation below, carried out from eq. (2.4) rather than by analogy, gives a different identity — linear rather than semilinear, and of the opposite sign relative to what was quoted — and has been checked three independent ways: symbolically against the general conformal-curvature-change formula, numerically against an arbitrary smooth test function, and against the closed-form textbook scale factor $\mu(\theta) = \sec^2(\theta/2)$ of the tangent stereographic projection, all in agreement. Theorem 5.2 below, and the maximum-principle bound of section 5.3, are both rebuilt on the corrected identity; neither survives unchanged from an earlier draft that used the original, incorrect one.

**Lemma 5.1** (The log-scale factor satisfies a linear Poisson equation for the sphere's own Laplace–Beltrami operator)**.** *Let $f$ be a conformal projection of a region $U$ of the sphere (radius $R$, metric $ds^2 = R^2 d\varphi^2 + R^2 \cos^2\varphi\, d\lambda^2$) with scale factor $\mu > 0$. Then*

$$\Delta_g(\ln \mu) \;=\; K \;=\; \frac{1}{R^2} \quad \text{on } U,$$

*where $\Delta_g$ is the Laplace–Beltrami operator of the sphere's own metric,*

$$\Delta_g h \;=\; \frac{1}{R^2 \cos\varphi}\left[\frac{\partial}{\partial\varphi}\left(\cos\varphi\,\frac{\partial h}{\partial\varphi}\right) + \frac{1}{\cos\varphi}\frac{\partial^2 h}{\partial\lambda^2}\right],$$

not *the flat coordinate Laplacian. This is a genuinely* linear *Poisson equation, with constant source $K$, in contrast to the semilinear form asserted in an earlier draft.*

*Proof.* Because $f$ is conformal, the target plane metric, pulled back through $f$ to the sphere's $(\varphi, \lambda)$-chart, is $\mu^2\big(R^2 d\varphi^2 + R^2\cos^2\varphi\, d\lambda^2\big)$, i.e. the metric with orthogonal coefficients $E = \mu^2 R^2$, $G = \mu^2 R^2 \cos^2\varphi$. This pulled-back metric is literally the flat Euclidean metric of the image plane, re-expressed in the $(\varphi, \lambda)$-chart via the diffeomorphism $f$, and curvature is a diffeomorphism invariant, so its Gaussian curvature computed from $(E, G)$ via eq. (2.4) must vanish identically:

$$0 \;=\; -\frac{1}{2\sqrt{EG}}\left[\frac{\partial}{\partial\varphi}\left(\frac{G_\varphi}{\sqrt{EG}}\right) + \frac{\partial}{\partial\lambda}\left(\frac{E_\lambda}{\sqrt{EG}}\right)\right].$$

Writing $E = e^{2\ln\mu}E_0$, $G = e^{2\ln\mu}G_0$ for the sphere's own coefficients $E_0 = R^2$, $G_0 = R^2\cos^2\varphi$, this is exactly the statement that the metric $e^{2\ln\mu}g_0$ (where $g_0$ is the sphere's own metric) is flat, and the standard conformal-change-of-curvature formula in two dimensions, $K_{e^{2\sigma}g_0} = e^{-2\sigma}(K_{g_0} - \Delta_{g_0}\sigma)$ (verified directly for this $(E_0, G_0)$ by explicit computation via eq. (2.4) rather than merely cited), gives, with $\sigma = \ln\mu$ and $K_{g_0} = K = 1/R^2$ (chapter 2):

$$0 \;=\; \mu^{-2}\left(K - \Delta_g(\ln\mu)\right) \quad\Longrightarrow\quad \Delta_g(\ln\mu) = K,$$

using $\mu \neq 0$ throughout ($\mu$ is a scale factor, nowhere zero) to divide through. This was checked directly, not only via the general conformal-change formula: substituting the tangent stereographic projection's own closed-form scale factor $\mu(\theta) = \sec^2(\theta/2)$ (colatitude $\theta = \pi/2 - \varphi$) into $\Delta_g(\ln \mu)$ gives $1/R^2$ identically, matching the claimed identity pointwise for this explicit, textbook case. □

Integrate lemma 5.1 over a region $U$ (a topological disk, for simplicity, with smooth boundary $\partial U$) and apply the Riemannian divergence theorem for the sphere's own metric, $\iint_U \Delta_g(\ln \mu)\, dA_g = \oint_{\partial U} \partial_n(\ln \mu)\, ds_g$, where $\partial_n$ is the outward normal derivative and $dA_g, ds_g$ are the sphere's own (source-surface) area and arclength elements:

$$\oint_{\partial U} \partial_n(\ln \mu)\, ds_g \;=\; \iint_U K\, dA_g \;=\; K \cdot \mathrm{Area}_g(U), \tag{5.1}$$

using $K \equiv 1/R^2$ constant to pull it out of the integral. This is already the complete, exact, and unconditional flux identity: unlike the semilinear version in an earlier draft, no further specialization (boundary normalization, or an additional hypothesis on the interior) is needed to reach a clean closed form, because the source term is constant. The total outward flux of $\ln \mu$ across $\partial U$, with respect to the sphere's own metric, is pinned down completely by $K$ and the *source* area of $U$ — not the image area, and not conditionally.

**Theorem 5.2** (Linear flux identity for the boundary log-scale factor)**.** *Let $U$ be a topological disk on the sphere with smooth boundary $\partial U$, and let $f$ be a conformal projection of $U$ with scale factor $\mu$. Then eq.* (5.1) *holds exactly: $\oint_{\partial U} \partial_n(\ln \mu)\, ds_g = K \cdot \mathrm{Area}_g(U)$, a fixed positive number determined purely by $U$'s own (source, spherical) area, independent of which conformal projection formula is chosen. Equivalently, combining with Gauss–Bonnet,*

$$\iint_U K\, dA_g \;+\; \oint_{\partial U} \kappa_g\, ds_g \;=\; 2\pi \tag{5.2}$$

*(where $\kappa_g$ is the geodesic curvature of $\partial U$), the same total flux equals $2\pi - \oint_{\partial U} \kappa_g\, ds_g$: the boundary's own geometry (via its geodesic curvature), not the choice of conformal projection, is what determines the total flux, with the projection formula free only to redistribute that fixed total around $\partial U$ and through the interior.*

*Proof.* Equation (5.1) is lemma 5.1 integrated via the Riemannian divergence theorem, as derived above. Equation (5.2) is the Gauss–Bonnet theorem for a topological disk, $\chi(U) = 1$, applied without modification (not re-derived here; see any standard reference, e.g. Do Carmo §4.5 — one of the few results in this volume, alongside theorem 2.1 itself, used as a cited black box rather than re-derived, a choice recorded explicitly rather than left implicit). Equating the two expressions for $\iint_U K\, dA_g$ gives the stated equivalence directly; no boundary normalization or interior hypothesis is needed anywhere in this argument, in contrast to the earlier, incorrect version of this theorem. □

**Key Result**

*Theorem 5.2 is unconditional and exact: the total sphere-metric flux of $\ln \mu$ out of $\partial U$ is fixed by $U$ alone (via either its own area or, equivalently, its boundary's geodesic curvature), with zero freedom for the projection designer at the level of the total, only at the level of how that fixed total is distributed. This is a strictly stronger and cleaner statement than the earlier, incorrect version: that version needed $\mu \equiv 1$ throughout $U$ (not just on $\partial U$) to reach a clean closed form and, even then, only approximately; the corrected identity needs no such hypothesis and holds exactly, precisely because the source term $K$ in lemma 5.1 is genuinely constant rather than depending on $\mu$.*

## 5.3 From flux to pointwise: a maximum-principle bound

Theorem 5.2 bounds a *flux integral* of $\ln \mu$, not $\ln \mu$ itself pointwise. Converting it into a genuine pointwise statement about where $\mu$ can be largest or smallest is possible directly from lemma 5.1, and — because the corrected equation is linear with constant source, rather than semilinear — the argument is if anything more standard than an earlier draft's (already elementary) attempt at this, and the conclusion points the opposite way: the corrected source term is *positive*, not negative, so the roles of maximum and minimum are exchanged relative to that earlier draft.

**Lemma 5.3** (Strong maximum principle for the log-scale factor)**.** *Let $U$ be a topological disk on the sphere with smooth boundary $\partial U$, as in theorem 5.2, and let $f$ be a conformal projection of $U$ with scale factor $\mu > 0$, $\mu \in C^2(\overline{U})$. Then $u := \ln \mu$ has no interior local maximum on $U$. Consequently*

$$\max_{\overline{U}} \mu \;=\; \max_{\partial U} \mu,$$

*and every maximizer of $\mu$ on $\overline{U}$ lies on $\partial U$; no interior point of $U$ can be a point of maximal scale.*

*Proof.* By lemma 5.1, $\Delta_g u = K$ at every point of $U$, and $K \equiv 1/R^2 > 0$ is the sphere's constant curvature, so $\Delta_g u(P) > 0$ strictly for every $P \in U$, with no point of degeneracy anywhere.

Suppose, for contradiction, that $u$ attains a local maximum at some interior point $P_0 \in U$. In a local orthonormal frame at $P_0$ for the sphere's own metric, the ordinary second-derivative test for a local maximum requires the (intrinsic, Riemannian) Hessian of $u$ at $P_0$ to be negative semi-definite; in particular its trace, which is exactly $\Delta_g u(P_0)$, must satisfy $\Delta_g u(P_0) \le 0$, directly contradicting $\Delta_g u(P_0) > 0$. Hence $u$ has no interior local maximum, and since exp is strictly increasing, neither does $\mu = e^u$.

$\overline{U} = U \cup \partial U$ is compact and $\mu$ continuous on it, so $\mu$ attains a global maximum somewhere on $\overline{U}$ (extreme value theorem); that maximizer cannot lie in $U$, by the previous paragraph, so it lies on $\partial U$, giving $\max_{\overline{U}} \mu = \max_{\partial U} \mu$. □

**Corollary 5.4** (Interior compression for boundary-true-to-scale conformal projections)**.** *If, in addition to the hypotheses of lemma 5.3, $\mu \equiv 1$ on $\partial U$ (the projection is exactly true-to-scale along the boundary curve, a normalization used by essentially every named conformal projection in* No Map Is Best *along at least one standard parallel or the central point), then $\mu(P) < 1$ for every interior point $P \in U$.*

*Proof.* By lemma 5.3, $\max_{\overline{U}} \mu = \max_{\partial U} \mu = 1$, attained only on $\partial U$. Every interior point $P$ therefore satisfies $\mu(P) \le 1$ (else it would exceed the global maximum), and since the maximum value 1 is attained only on $\partial U$, $\mu(P) \ne 1$ for interior $P$. Together these give $\mu(P) < 1$. □

**Key Result**

*Lemma 5.3 and corollary 5.4 give the pointwise statement section 5.3 sets out to find, and — now that the underlying equation is known to be linear rather than semilinear — the argument needs nothing beyond the ordinary second-derivative test, exactly as an elementary maximum-principle argument should. The direction is worth stating plainly since it reverses the (incorrect) earlier draft's conclusion: a conformal projection of a spherical region can never be maximally* expanded*, in the pointwise scale sense, anywhere in the interior of that region — the point or points of greatest linear scale are always on $\partial U$, never inside. Sharpened for the boundary-true-to-scale case, the interior is strictly* compressive*, $\mu < 1$ at literally every interior point: a conformal projection cannot be true-to-scale on its boundary and true-to-scale, let alone expansive, at even a single interior point. This matches direct inspection of a concrete case — a secant-type polar stereographic projection with its single standard parallel placed exactly on $\partial U$ has $\mu(\theta) = \sec^2(\theta/2)/\sec^2(\alpha/2)$ for boundary colatitude $\alpha$, strictly less than 1 for every colatitude $\theta < \alpha$ — which is the same concrete case used to check lemma 5.1 itself in the first place.*

This box's quantitative request — an explicit bound on $1 - \min_U \mu$ in the style of theorem 5.10's $\delta_0(\alpha)$ — is answered in full, and exactly rather than merely as an inequality, by theorem 5.6 immediately below: the boundary-normalized geodesic-disk case admits the closed form $\mu(\theta) = \cos^2(\alpha/2)/\cos^2(\theta/2)$, giving $\min_{\overline{D_\alpha}} \mu = \cos^2(\alpha/2)$ exactly, attained uniquely at the centre. The comparison against $\delta_0(\alpha)$ this box asked for is carried out in the keyresult following that theorem's proof (leading-order gap $\frac{1}{6}\alpha^2$), and the further generalization to an arbitrary, non-rotationally-symmetric boundary profile is carried out afterward in theorem 5.8. Nothing in this closure required Milnor's own text; it follows entirely from the corrected linear Poisson equation already established earlier in this chapter, via the Green's-function/separation-of-variables route this box anticipated. The separate, still-open question of Milnor's own stated relationship (if any) between $\delta_0(\alpha)$ and total curvature does require his primary 1969 text and is tracked by its own todobox in section 5.5 and section 5.6 below.

## 5.4 The quantitative bound: an exact solution on a geodesic disk

The todobox above asks for a genuine number, not only a qualitative direction. For the boundary-normalized case singled out throughout this chapter ($\mu \equiv 1$ on $\partial U$) and for $U$ a geodesic disk specifically, the linear Poisson equation of lemma 5.1 turns out to admit something stronger than a barrier-comparison bound: a closed-form, exact solution, obtained by an argument in two independent steps — first that the solution must be rotationally symmetric, then solving the resulting ordinary differential equation directly.

**Lemma 5.5** (Rotational symmetry of the solution)**.** *Let $D_\alpha$ be the geodesic disk of angular radius $\alpha$ centred at a point $O$ of the sphere, and let $\mu > 0$ be any $C^2$ function on $\overline{D_\alpha}$ satisfying $\Delta_g(\ln \mu) = K$ on $D_\alpha$ (lemma 5.1) and $\mu \equiv 1$ on $\partial D_\alpha$. Then $\ln \mu$ depends only on the geodesic distance $\theta$ from $O$ (equivalently, on colatitude $\theta$ in a polar coordinate system centred at $O$), not on the azimuthal coordinate.*

*Proof.* Fix any rotation $R_\psi$ of the sphere about the axis through $O$, by angle $\psi$. Both $D_\alpha$ and the operator $\Delta_g$ are invariant under $R_\psi$, since $R_\psi$ is an isometry of the sphere fixing $O$ and $D_\alpha$ is exactly the set of points within geodesic distance $\alpha$ of $O$. Consequently $v := \ln \mu \circ R_\psi$ satisfies the same equation, $\Delta_g v = K$ on $D_\alpha$, and the same boundary condition, $v \equiv 0$ on $\partial D_\alpha$ (since $R_\psi$ maps $\partial D_\alpha$ to itself). The Dirichlet problem for the linear Poisson equation $\Delta_g w = K$ on a bounded domain, with $K$ and the boundary data fixed, has at most one $C^2(\overline{D_\alpha})$ solution: if $w_1, w_2$ both solve it, $w_1 - w_2$ is $\Delta_g$-harmonic on $D_\alpha$ with zero boundary values, hence identically zero by the maximum principle applied to both $w_1 - w_2$ and $w_2 - w_1$

(lemma 5.3's argument, run on a harmonic rather than superharmonic function, forbids both an interior maximum and an interior minimum unless the function is constant). Applying this uniqueness to $v = \ln\mu \circ R_\psi$ and $\ln\mu$ itself, which solve the identical problem, gives $\ln\mu \circ R_\psi = \ln\mu$ for every $\psi$: $\ln\mu$ is invariant under every rotation about $O$, i.e. depends only on $\theta$. □

**Key Result**

*Lemma 5.5 is worth pausing on before proceeding to solve for $\mu(\theta)$ explicitly, because its content is stronger than it may first appear: it says that* every *conformal projection of a geodesic disk that is true-to-scale along the entire boundary circle — azimuthal or not, however it is constructed — has exactly the same scale-factor profile as a function of distance from the centre. The freedom a projection designer has, once the boundary normalization is fixed, is confined entirely to how $\mu$'s (already pinned-down) values are distributed* angularly *at each radius, if the map is allowed to be non-azimuthal; the radial profile itself has no remaining freedom at all.*

With lemma 5.5 in hand, write $u(\theta) := \ln\mu(\theta)$ and use the standard form of the Laplace–Beltrami operator on the sphere (radius $R$) acting on a function of colatitude $\theta$ alone,

$$\Delta_g u = \frac{1}{R^2 \sin\theta}\frac{d}{d\theta}\left(\sin\theta\,\frac{du}{d\theta}\right),$$

so that lemma 5.1's equation $\Delta_g u = K = 1/R^2$ becomes the ordinary differential equation $\dfrac{d}{d\theta}\left(\sin\theta\,\dfrac{du}{d\theta}\right) = \sin\theta$.

**Theorem 5.6** (Exact scale factor on a boundary-normalized geodesic disk)**.** *Let $D_\alpha$ be a geodesic disk of angular radius $\alpha$ on the sphere of radius $R$, and let $f$ be any conformal projection of $D_\alpha$ with scale factor $\mu > 0$ satisfying $\mu \equiv 1$ on $\partial D_\alpha$. Then, for every point of $D_\alpha$ at geodesic colatitude $\theta \in [0, \alpha]$ from the centre,*

$$\mu(\theta) \;=\; \frac{\cos^2(\alpha/2)}{\cos^2(\theta/2)}. \tag{5.3}$$

*In particular $\mu$ is strictly increasing in $\theta$ on $[0, \alpha]$, with*

$$\min_{\overline{D_\alpha}}\mu \;=\; \mu(0) \;=\; \cos^2(\alpha/2), \qquad \max_{\overline{D_\alpha}}\mu \;=\; \mu(\alpha) \;=\; 1,$$

*attained uniquely at the centre and uniformly on the boundary respectively.*

*Proof.* By lemma 5.5, seek $u(\theta)$ solving $\frac{d}{d\theta}(\sin\theta\,u') = \sin\theta$ with $u(\alpha) = 0$ and the regularity condition forced by smoothness of $\mu$ at the centre $O$ of the disk (a point where the polar coordinate $\theta$ is singular): $\sin\theta\,u'(\theta) \to 0$ as $\theta \to 0^+$. Integrating once, $\sin\theta\,u'(\theta) = -\cos\theta + C_1$ for a constant $C_1$; the regularity condition at $\theta = 0$ forces $C_1 = 1$ (else $\sin\theta\,u'(\theta) \to C_1 - 1 \neq 0$), giving

$$u'(\theta) = \frac{1-\cos\theta}{\sin\theta} = \tan(\theta/2)$$

by the half-angle identity. Integrating again, $u(\theta) = -2\ln\cos(\theta/2) + C_2$, and the boundary condition $u(\alpha) = 0$ fixes $C_2 = 2\ln\cos(\alpha/2)$, giving $u(\theta) = 2\ln\big(\cos(\alpha/2)/\cos(\theta/2)\big)$ and hence eq. (5.3) on exponentiating. Monotonicity follows from $d\mu/d\theta = 2\mu(\theta)\tan(\theta/2) \geq 0$ on $[0, \alpha]$, with equality only at $\theta = 0$, and the stated extremes follow directly by evaluation at the endpoints, consistent with (and now made numerically exact by) the qualitative ordering already established in lemma 5.3 and corollary 5.4. □

**Epistemic Notice**

[Three independent checks on eq. (5.3)] This closed form was verified three ways before being written into this section, in keeping with the working method of chapter 1. First, direct substitution of eq. (5.3) into the coordinate form of $\Delta_g$ above, carried out symbolically, returns exactly $1/R^2$ with zero residual. Second, it coincides exactly, not merely approximately, with the concrete polar-stereographic formula $\mu(\theta) = \sec^2(\theta/2)/\sec^2(\alpha/2)$ already used as a worked example in the proof of corollary 5.4 — itself independently checked there against the textbook tangent-stereographic scale formula. Third, and most independently of the symbolic derivation above, the boundary-value problem was solved numerically from scratch (a standard collocation BVP solver applied directly to the ODE $\frac{d}{d\theta}(\sin\theta\, u') = \sin\theta$ with the stated boundary and regularity conditions, without using the closed-form antiderivatives at any step) for a representative value $\alpha = 0.9\,\text{rad}$; the numerical solution agreed with eq. (5.3) to within $10^{-12}$ at nine test points spanning $[0, \alpha]$, which is consistent with solver tolerance and not with a systematic discrepancy.

**Key Result**

*Theorem 5.6 answers the todobox's quantitative question completely for the boundary-normalized geodesic-disk case, and does so with an exact identity rather than an inequality: the minimum scale anywhere in $D_\alpha$ is not merely bounded below but pinned down exactly at $\cos^2(\alpha/2)$, attained only at the centre. Writing the corresponding conformal-restricted, boundary-normalized analogue of Milnor's distortion as $\delta_{\text{conf}}(\alpha) := \ln\big(\max\mu/\min\mu\big) = -2\ln\cos(\alpha/2)$ and comparing small-angle expansions against eq. (5.7): $\delta_{\text{conf}}(\alpha) = \frac{1}{4}\alpha^2 + O(\alpha^4)$, versus Milnor's $\delta_0(\alpha) = \frac{1}{6}\alpha^2 + O(\alpha^4)$. The conformal, boundary-normalized quantity is strictly larger to leading order ($\frac{1}{4} > \frac{1}{6}$), exactly as it must be: $\delta_0(\alpha)$ is Milnor's infimum over* all *homeomorphisms with* no *boundary normalization imposed, so no conformal, boundary-constrained construction can beat it, and theorem 5.6 shows the specific gap is not merely present but has an exact, computable leading-order size, $\frac{1}{6}\alpha^2$ short of Milnor's optimum, itself a small but genuine quantitative consequence of adding the boundary constraint on top of conformality.*

## 5.5 The general boundary-value problem: an explicit Poisson kernel

Theorem 5.6 settles the case $\mu \equiv 1$ on $\partial D_\alpha$. The natural next question — left open at the end of the previous section — is what happens for an arbitrary boundary profile, $\mu|_{\partial D_\alpha} = b(\lambda) > 0$ a general (not necessarily constant, not necessarily rotationally symmetric) function of the azimuthal coordinate $\lambda$ alone. This turns out to be fully solvable in closed form as well, by transporting the problem to the flat plane through the tangent stereographic projection and invoking a classical fact about the Laplace operator in two dimensions.

**Lemma 5.7** (Conformal invariance of harmonicity in two dimensions)**.** *Let $g_1 = e^{2\sigma} g_2$ be two Riemannian metrics on the same two-dimensional coordinate chart, related pointwise by a conformal factor $e^{2\sigma}$. Then for any $C^2$ function $h$ on the chart, $\Delta_{g_1} h = e^{-2\sigma}\,\Delta_{g_2} h$. In particular $\Delta_{g_1} h = 0$ if and only if $\Delta_{g_2} h = 0$: harmonicity is a conformally invariant condition in two (and only two) dimensions.*

*Proof.* This is the standard formula for the Laplace–Beltrami operator under a conformal rescaling of a two-dimensional metric, and is most transparent in isothermal coordinates, where any 2D metric can locally be written $g = e^{2\tau}(dx^2 + dy^2)$ for some $\tau$; then $\Delta_g h = e^{-2\tau}(h_{xx} + h_{yy})$ by direct computation from the general orthogonal-coordinates curvature-and-Laplacian formulas of chapter 2 (the same machinery underlying eq. (2.4)), applied with $E = G = e^{2\tau}$, $F = 0$. Writing $g_2 = e^{2\tau}(dx^2 + dy^2)$ and $g_1 = e^{2\sigma} g_2 = e^{2(\sigma+\tau)}(dx^2 + dy^2)$ in the same isothermal coordinates gives $\Delta_{g_2} h = e^{-2\tau}(h_{xx} + h_{yy})$ and $\Delta_{g_1} h = e^{-2(\sigma+\tau)}(h_{xx} + h_{yy}) = e^{-2\sigma}\Delta_{g_2} h$ directly. The restriction to exactly two dimensions is essential: in $d \geq 3$ the

analogous formula acquires an additional zero-order term in $h$ and $\Delta_{g_1} h = 0$ no longer follows from $\Delta_{g_2} h = 0$ alone. □

**Key Result**

*Lemma 5.7 is the reason the corrected identity of lemma 5.1 could be derived at all: the curvature-change formula used there, $K_{e^{2\sigma} g_0} = e^{-2\sigma}(K_{g_0} - \Delta_{g_0}\sigma)$, and the present Laplacian-change formula are companion facts about the same conformal rescaling, one for curvature and one for the Laplacian, and both are special to two dimensions. Here it is the Laplacian statement, not the curvature one, that is needed: it says a function's harmonicity does not depend on* which *metric, among a conformal family sharing the same coordinate chart, is used to test it.*

Apply lemma 5.7 with $g_1$ the pullback to the sphere's $(\theta, \lambda)$-chart of the flat metric of the plane under the tangent stereographic projection from the point antipodal to $D_\alpha$'s centre, and $g_2$ the sphere's own metric $g$: by the defining property of the scale factor (chapter 3), $g_1 = \mu_{\mathrm{st}}^2\, g$ in this shared chart, where $\mu_{\mathrm{st}}(\theta) = \sec^2(\theta/2)$ is the tangent stereographic projection's own scale factor, already used repeatedly in chapter 5. Concretely, writing $r = 2R\tan(\theta/2)$ for the image radius under this projection (the standard tangent-stereographic radial formula, consistent with $\mu_{\mathrm{st}} = \sec^2(\theta/2)$), a function $h(r, \lambda)$ satisfying the ordinary flat-plane Laplace equation $\frac{1}{r}\partial_r(r\,\partial_r h) + \frac{1}{r^2}\partial_\lambda^2 h = 0$ pulls back, via $r = 2R\tan(\theta/2)$, to a function of $(\theta, \lambda)$ satisfying $\Delta_g h = 0$ on $D_\alpha$ — $\Delta_g$-harmonic on the sphere, exactly as claimed by lemma 5.7.

**Theorem 5.8** (General closed-form solution on a geodesic disk)**.** *Let $D_\alpha$ be a geodesic disk of angular radius $\alpha$ on the sphere of radius $R$, let $\rho_0 := 2R\tan(\alpha/2)$, and let $b : [0, 2\pi) \to (0, \infty)$ be any continuous boundary profile, $\mu|_{\partial D_\alpha}(\lambda) = b(\lambda)$. Then the unique $C^2(\overline{D_\alpha})$ solution of $\Delta_g(\ln\mu) = K$ on $D_\alpha$ with this boundary data is*

$$\ln\mu(\theta,\lambda) \;=\; 2\ln\frac{\cos(\alpha/2)}{\cos(\theta/2)} \;+\; \frac{1}{2\pi}\int_0^{2\pi} \ln b(\lambda')\,\frac{\rho_0^2 - r^2}{\rho_0^2 - 2\rho_0 r\cos(\lambda-\lambda') + r^2}\,d\lambda', \qquad r = 2R\tan(\theta/2). \tag{5.4}$$

*Proof.* Uniqueness is lemma 5.5's proof verbatim, which used only that the equation is linear with the given fixed data, not that the data is rotationally symmetric — the Dirichlet problem for $\Delta_g w = K$ with fixed boundary values has at most one $C^2(\overline{D_\alpha})$ solution regardless of whether those values are constant. For existence, write $u = \ln\mu$ as the sum of the particular solution $u_p(\theta) := 2\ln(\cos(\alpha/2)/\cos(\theta/2))$ of theorem 5.6 (which satisfies $\Delta_g u_p = K$ and $u_p|_{\partial D_\alpha} = 0$) and a correction $h$ with $\Delta_g h = 0$ on $D_\alpha$ and $h|_{\partial D_\alpha} = \ln b(\lambda) - 0 = \ln b(\lambda)$; then $u = u_p + h$ satisfies $\Delta_g u = K + 0 = K$ by linearity and $u|_{\partial D_\alpha} = 0 + \ln b(\lambda) = \ln b(\lambda)$, as required. For $h$, pull back the classical Poisson kernel solution of the flat Dirichlet problem on the disk of radius $\rho_0$ in the plane, $h_0(r,\lambda) = \frac{1}{2\pi}\int_0^{2\pi} \ln b(\lambda')\,\frac{\rho_0^2 - r^2}{\rho_0^2 - 2\rho_0 r\cos(\lambda-\lambda') + r^2}\,d\lambda'$ (standard 2D potential theory, not re-derived here), which is flat-harmonic on the open disk and attains the boundary values $\ln b(\lambda)$ continuously as $r \to \rho_0^-$; by lemma 5.7 and the discussion above, $h(\theta,\lambda) := h_0(2R\tan(\theta/2), \lambda)$ is $\Delta_g$-harmonic on $D_\alpha$ with the same boundary values. Adding gives eq. (5.4). □

**Epistemic Notice**

[Verification of theorem 5.8] Three checks were run, none of them merely a restatement of the derivation. First, the key structural claim underlying lemma 5.7's application — that flat-harmonic functions on the image disk pull back, via $r = 2R\tan(\theta/2)$, to $\Delta_g$-harmonic functions on $D_\alpha$ — was checked symbolically against the elementary flat-harmonic basis functions $r^n\cos(n\lambda)$ for $n = 1, 2, 3$, substituting $r = 2R\tan(\theta/2)$ into the full two-variable coordinate expression for $\Delta_g$ and obtaining exactly zero in every case (this basis spans, together with $r^n\sin(n\lambda)$ and constants, all bounded harmonic functions on the flat disk via Fourier series, so the check is representative of the general claim rather than an isolated instance of it). Second, for a concrete non-constant, non-symmetric boundary profile $\ln b(\lambda) = 0.2 + 0.3\cos\lambda - 0.1\sin 3\lambda$ and $\alpha = 0.9$, eq. (5.4) was evaluated numerically at an interior point away from the centre ($\theta = 0.5$, $\lambda = 1.2$) and $\Delta_g u$ was computed independently by finite differences on $(\theta, \lambda)$; the result matched $K = 1/R^2$ to six decimal places. Third, evaluating eq. (5.4) at points approaching the boundary ($r/\rho_0 \to 1^-$) at a fixed $\lambda$ was checked to converge to the prescribed boundary value $\ln b(\lambda)$, confirming the boundary condition is genuinely recovered and not merely satisfied by construction of the kernel.

**Key Result**

*[The centre value depends on the boundary only through its average] Setting $\theta = 0$ (so $r = 0$) in eq. (5.4), the kernel collapses to the constant $\rho_0^2/\rho_0^2 = 1$, giving the mean-value identity*

$$\ln\mu(O) \;=\; 2\ln\cos(\alpha/2) \;+\; \frac{1}{2\pi}\int_0^{2\pi} \ln b(\lambda')\,d\lambda', \tag{5.5}$$

*i.e. the log-scale factor at the disk's centre is the boundary-normalized value $2\ln\cos(\alpha/2)$ of theorem 5.6* plus *the plain average of $\ln b$ over the boundary circle, with no dependence whatsoever on how $b$ is distributed around that circle beyond its mean. This was checked numerically for the same test profile as above: the predicted value $2\ln\cos(\alpha/2) + \overline{\ln b}$ and the direct numerical evaluation of eq. (5.4) at $\theta \to 0$ agreed to five decimal places. It is a genuine strengthening of theorem 5.6's mean-value-free case: no matter how unevenly a boundary-normalized projection's true-to-scale condition is relaxed around $\partial D_\alpha$ — concentrated in one arc, oscillating, arbitrary — the resulting scale factor at the very centre of the disk is governed by one number alone, the boundary average, and theorem 5.6 is recovered exactly as the special case $b \equiv 1$.*

**Todo — verify primary source**

Theorem 5.8 closes item (i) of the previous todobox in full generality (any continuous boundary profile, not only the rotationally symmetric ones lemma 5.5's method alone would reach), leaving a single item open in this chapter: whether eq. (5.3), eq. (5.4), or the gap $\delta_0(\alpha) - \delta_{\text{conf}}(\alpha)$ identified after theorem 5.6, connects to item (c) of the todobox in section 5.6 below (Milnor's own stated relationship, if any, between $\delta_0(\alpha)$ and total curvature), which requires checking against Milnor's primary 1969 text rather than further computation from material already established in this volume.

## 5.6 Milnor's theorem

John Milnor's 1969 note addresses a version of exactly the question flagged as open at the end of section 5.2: not a flux identity, but a genuine bound on the *worst-case* scale distortion over an entire region, for a specific and cartographically natural region — a geodesic disk on the sphere of given angular radius — and for a global (not merely infinitesimal) distortion measure defined directly from geodesic distances, so that the result applies to arbitrary homeomorphisms, not only to the smooth, locally-linearizable projections to which theorem 3.1's indicatrix construction applies.

**Epistemic Notice**

The publication record for Milnor's paper — *A Problem in Cartography*, *American Mathematical Monthly* 76(10), pp. 1101–1112, published December 1969 — has now been confirmed directly against the journal's own listing and against the bibliography of Milnor's collected papers. The statement of the theorem given below, including the exact defining formula for Milnor's distortion functional $\delta$, the sharp bound $\delta_0(\alpha)$, its small-angle asymptotics, and the identification of the extremal projection, was originally drawn from a single secondary exposition (a 2023 applied-mathematics blog post by J. D. Cook). It has since been independently corroborated by a second, substantially stronger secondary source: Papadopoulos, *Quasiconformal mappings, from Ptolemy's geography to the work of Teichmüller* (arXiv:1702.03756; final version in *Uniformization, Riemann–Hilbert Correspondence, Calabi–Yau Manifolds, and Picard–Fuchs Equations*, ed. Ji & Yau, International Press, 2017), a peer-reviewed historical survey that reports Milnor's own definitions and cites explicitly "He obtains a lower bound for this ratio in the case of the so-called azimuthal equidistant projection, and he proves the existence and uniqueness of a minimum distortion map in this case," matching the blog's account of theorem 5.10 independently and by an author who states he is summarizing Milnor's actual exposition, not merely citing the paper's existence. Two secondary sources with no evident relationship to one another agreeing on the same formula, uniqueness claim, and extremal projection meaningfully raises confidence, but neither is a direct reading of Milnor's own proof, and the derivation below is reconstructed from the stated result, not reproduced from Milnor's argument; this distinction is maintained explicitly in the todobox that follows.

**Definition 5.9** (Milnor's global distortion functional)**.** Let $f$ be a homeomorphism from a region $U$ of the sphere (with intrinsic geodesic distance $d_S$) onto a region of the Euclidean plane (with distance $d_E$). For any two distinct points $x, y \in U$, the *scale* of $f$ with respect to the pair is $d_E(f(x), f(y))/d_S(x, y)$. Let

$$\sigma_1 = \inf_{x \neq y \in U} \frac{d_E(f(x), f(y))}{d_S(x, y)}, \qquad \sigma_2 = \sup_{x \neq y \in U} \frac{d_E(f(x), f(y))}{d_S(x, y)},$$

the minimum and maximum scale attained anywhere on $U$. Milnor's *distortion* of $f$ is

$$\delta(f) = \log(\sigma_2/\sigma_1). \tag{5.6}$$

Note immediately that $\delta(f)$ is a strictly global quantity, built from the extremes of the scale ratio over *all pairs* of points in $U$, not from the pointwise principal scales $(a, b)$ of theorem 3.1 at a single point; for a smooth projection, $\sigma_1$ and $\sigma_2$ will typically (though this requires an argument not given here) coincide with the infimum and supremum, over $U$, of the pointwise minimum and maximum principal scales $b(P)$ and $a(P)$ respectively, which is the sense in which definition 5.9 generalizes, rather than replaces, the local machinery of chapter 3.

**Theorem 5.10** (Milnor, 1969)**.** *Let $D_\alpha$ be the geodesic disk of angular radius $\alpha$ on the sphere, $0 < \alpha < \pi$. Among all homeomorphisms $f$ of $D_\alpha$ onto a region of the plane, the minimum possible distortion is*

$$\delta_0(\alpha) = \log\Big(\frac{\alpha}{\sin \alpha}\Big), \tag{5.7}$$

*and this minimum is attained, uniquely up to similarity transformations (translation, rotation, and uniform rescaling) of the plane, by the* azimuthal equidistant projection *centred at the disk's centre.*

*Proof.* Not reproduced here; see the todobox below. □

**Key Result**

*Equation* (5.7) *makes the qualitative expectation of section 5.2 completely quantitative and, unlike theorem 5.2, gives an honest* pointwise-strength *bound rather than only a flux identity: no map of a spherical cap of angular radius $\alpha$, by any construction whatsoever (not merely conformal ones), can bring the ratio of its best-case to worst-case scale below $\alpha/\sin\alpha$. Two checks confirm the bound behaves as it must. As $\alpha \to 0$, $\alpha/\sin\alpha \to 1$ and $\delta_0(\alpha) \to 0$: an arbitrarily small cap can be mapped with arbitrarily little distortion, consistent with the sphere's local flatness and with corollary 2.2 only forbidding* exact*, not approximate, isometry. The small-angle expansion $\sin\alpha = \alpha - \alpha^3/6 + O(\alpha^5)$ gives $\alpha/\sin\alpha = 1 + \alpha^2/6 + O(\alpha^4)$, so $\delta_0(\alpha) \approx \alpha^2/6$ for small $\alpha$; since the area of $D_\alpha$ on a unit sphere is $2\pi(1-\cos\alpha) \approx \pi\alpha^2$ and the total sphere's area is $4\pi$, the area fraction is $\approx \alpha^2/4$, so $\delta_0(\alpha) \approx \frac{2}{3}\cdot$(area fraction) to leading order: minimum distortion grows, to leading order, linearly in the fraction of the sphere's area covered, with a fixed constant $2/3$, not merely "some" unspecified growth. For a region the size of metropolitan France (area fraction of order $10^{-3}$ of the globe), this predicts a minimum achievable Milnor-distortion on the order of $7 \times 10^{-4}$: a genuinely tiny, but strictly positive and computable, unavoidable floor.*

*Remark* 5.11 (Why the extremal projection is azimuthal equidistant, plausibly). It is worth checking that theorem 5.10's extremal projection is at least plausible independently of the theorem's proof, since the azimuthal equidistant projection is already a familiar object: by construction it sends every point at geodesic distance $\rho$ from the centre to a point at Euclidean distance exactly $\rho$ from the image of the centre (in the natural normalization), so it is exactly scale-correct, $\sigma = 1$, along every radius through the centre; all of its distortion is confined to the transverse (angular) direction, where the scale factor at geodesic radius $\rho$ works out to $\rho/\sin\rho$ in the appropriate units, increasing monotonically from 1 at the centre to $\alpha/\sin\alpha$ at the disk's boundary $\rho = \alpha$. This makes $\sigma_1 = 1$ (attained at and near the centre) and $\sigma_2 = \alpha/\sin\alpha$ (attained at the boundary) plausible by direct inspection of the formula, giving $\delta = \log(\alpha/\sin\alpha)$ exactly as theorem 5.10 claims for this specific candidate; what is *not* independently checked here is the harder half of the theorem, that no other homeomorphism does better, which is the actual content of Milnor's proof and is exactly what remains to be verified against the primary source.

## 5.7 Resolving item (d): Milnor's paper contains two theorems, not one

The Papadopoulos survey cited above settles the question, left open in the previous draft's todobox, of how theorem 5.10 relates to the "Chebyshev principle" attributed to Milnor by Bermejo & Otero. It reports that Milnor's 1969 paper contains, *in addition to* the geodesic-disk theorem stated as theorem 5.10 above, a full exposition with proof of an earlier, different result of Chebyshev's, restricted to *conformal* maps:

> [Chebyshev's theorem, as reported in Milnor's exposition] Let $\Omega$ be a simply connected open region of the sphere with twice-differentiable boundary, and $f : \Omega \to \mathbb{E}$ a conformal map, with infinitesimal scale $\sigma(x) = \lim_{y\to x} d_E(f(x), f(y))/d_S(x, y)$ (well defined at each point precisely because $f$ is conformal). Among all such $f$, there is one, unique up to similarity, minimizing $\sup_\Omega \sigma/\inf_\Omega \sigma$, and it is characterized by $\sigma$ being constant on $\partial\Omega$.

**Key Result**

*This is a genuinely different theorem from theorem 5.10, on three counts, not one: (1) it is restricted to* conformal *$f$, whereas theorem 5.10 ranges over* all *homeomorphisms; (2) its distortion measure is a ratio of* infinitesimal *scales,* $\sup\sigma/\inf\sigma$*, whereas definition 5.9's* $\delta(f) = \log(\sigma_2/\sigma_1)$ *is built from* finite*, whole-region point-pair distances; (3) it applies to an arbitrary twice-differentiable-boundary region* $\Omega$*, characterized by a boundary condition, whereas theorem 5.10 is specific to a geodesic disk, with no boundary condition needed because the extremal azimuthal equidistant projection is pinned down by symmetry instead. Bermejo & Otero's description — distortion as the oscillation of the log of infinitesimal scale, tied to a Chebyshev extremal principle — is this first theorem, not theorem 5.10; the two coexist in the same 1969 paper as, respectively, Milnor's own exposition of a result predating him (Chebyshev's, worked out with full proofs "for the first time, per the survey) and Milnor's own original contribution (the geodesic-disk bound). Todobox item (d) is accordingly resolved: the relationship is not a conflation to be avoided but two distinct, coexisting theorems correctly kept separate in definition 5.9–theorem 5.10 above, now traced to their right names.*

**Epistemic Notice**

This resolution rests on the Papadopoulos survey's own summary of Milnor's exposition, itself a secondary (though peer-reviewed and directly on-point) account, not on a direct reading of either Chebyshev's or Milnor's original arguments. It is nonetheless a materially stronger basis than the previous draft had for item (d), and is independent of the Cook blog post used for theorem 5.10 itself: the two secondary sources cover different parts of Milnor's paper (Cook: the geodesic-disk bound; Papadopoulos: the Chebyshev exposition) and corroborate rather than duplicate one another.

**Todo — verify primary source**

Theorem 5.10 is now stated with a specific, checkable formula and a named extremal projection, corroborated by two independent secondary sources, but the following remain to be completed against Milnor's primary 1969 text before the theorem can be considered fully verified for this volume's standards: (a) the exact statement and technique of the optimality proof (i.e. why *no* homeomorphism, smooth or not, can beat the azimuthal equidistant construction — the remark above only checks that the extremal candidate achieves the claimed value, not that nothing does better); (b) whether the uniqueness claim ("unique up to similarity") is exactly as Milnor stated it, including whatever regularity hypotheses on $f$ (continuity alone, or more) his argument actually requires; (c) the relationship, if any, Milnor's paper draws explicitly between $\delta_0(\alpha)$ and the total curvature of $D_\alpha$, which would connect theorem 5.10 directly to the Gauss–Bonnet bound of theorem 5.2 rather than leaving the two as parallel but logically unconnected results, as they are left in section 5.9 below. Item (d) of the previous version of this todobox — the relationship to the Bermejo–Otero Chebyshev principle — is resolved above in section 5.7 and removed from this list. Section 5.8 below substantially narrows item (b) and, incidentally, item (a) as it applies to Milnor's Chebyshev exposition specifically, though not as it applies to Milnor's own geodesic-disk theorem.

## 5.8 A closer reading of the Papadopoulos survey: Milnor's Chebyshev theorem in full, and a historical aside in Milnor's own words

The Papadopoulos survey already cited for section 5.7 was read this pass at a level of detail beyond what that section drew on: rather than the survey's summary claim that Milnor's paper reproduces a Chebyshev result with full proof, the survey's own restatement of that proof's conclusion — not the derivation, but the precise hypotheses and the precise assertion — was read directly.

**Key Result**

*[Milnor's own statement of the Chebyshev theorem, at second hand but in full] Papadopoulos's account gives what he presents as Milnor's own exposition, restricted to conformal maps $f$ from a simply connected open subset $\Omega$ of the sphere onto the Euclidean plane, of a theorem he attributes jointly to Chebyshev and to Milnor's own "revision" of the argument: if $\Omega$'s boundary is a twice-differentiable curve, there exists exactly one conformal map from $\Omega$ onto the plane, unique up to similarity, whose infinitesimal scale function is minimal in the sense of definition 5.9's ratio $\sup\sigma/\inf\sigma$; and this extremal map is characterized, not merely constructed, by the property that its infinitesimal scale is constant on the boundary of $\Omega$. This directly closes item (b) of the todobox above for the Chebyshev theorem specifically (section 5.7's "first theorem"): the uniqueness is up to similarity, exactly as this volume's own account already stated by inference, and the regularity hypothesis is a twice-differentiable boundary curve, not mere continuity, now confirmed rather than assumed. It remains explicitly unresolved for theorem 5.10 itself (Milnor's own geodesic-disk result, a different theorem by section 5.7's own three-count distinction), which Papadopoulos's survey treats far more briefly and without reproducing Milnor's regularity hypotheses in comparable detail.*

Two further details, incidental to section 5.8 but worth recording, sharpen this volume's picture of Milnor's paper without belonging to either open item above. First, Papadopoulos reports that Milnor's own exposition includes a wry aside on the Chebyshev theorem's practical fate: in Milnor's own words, as quoted by Papadopoulos, the result "has been available for more than a hundred years, but to my knowledge it has never been used by actual map makers." Second, and more substantively, Papadopoulos notes that the French mathematician Gaston Darboux independently supplied his own proof of the same Chebyshev theorem, by a different route through potential theory, and that Milnor's paper shows no sign of awareness of Darboux's work. Neither detail changes any claim already made in this chapter, but the second is worth flagging for future primary-source work: a direct reading of Milnor's 1969 text, when undertaken, should check this apparent gap in Milnor's own citations against the paper's actual bibliography, since a secondary survey's inference that an author was unaware of prior work is weaker evidence than the primary text's own reference list would be.

**Todo — verify primary source**

**Key Result**

*[A second, independent source: Darboux named and dated precisely, and the shared proof technique identified] A further pass located a second Papadopoulos paper, independent of the survey section 5.7 already cites:* Euler and Chebyshev: From the sphere to the plane and backwards *(arXiv:1608.02724, 2016), a different, earlier historical article by the same author, covering the same Chebyshev–Darboux–Milnor lineage from the classical-cartography side rather than the quasiconformal-theory side. It corroborates section 5.8's Darboux remark independently rather than merely repeating it, and adds two specifics neither the blog account nor the 2017 survey supplied: first, an exact citation for Darboux's paper — G. Darboux, "Sur la construction des cartes géographiques,"* Bulletin des sciences mathématiques *35 (1911), 23–28 — where before this volume had only "Darboux, by a different route through potential theory," unattributed to any specific publication; second, a precise statement, not merely of* that *Milnor proved Chebyshev's theorem but of the shared* technique*: both Darboux's 1911 proof and Milnor's 1969 proof (located by this source to a specific page, Milnor's p. 1111) proceed "based on the same idea, the solution of the Laplace equation," i.e. a Dirichlet-problem argument in which the boundary condition is the required constancy of the scale factor and the extremal map's scale function in the interior is the harmonic extension of that boundary data. This source states independently, in language matching but not copied from the 2017 survey, that "it seems that Milnor was not aware of the work of Darboux on the question," the same historical claim now resting on two independently authored papers rather than one.*

**Todo — verify primary source**

Section 5.8 narrows item (a) of section 5.6's todobox one further step, for the Chebyshev-theorem portion specifically: the proof *technique* is now identified by name (Dirichlet problem for Laplace's equation, boundary data fixed by the required constant-scale condition) and corroborated across two independent secondary sources rather than asserted from one, and Darboux's own paper is now precisely enough cited (*Bulletin des sciences mathématiques* 35, 1911, pp. 23–28) that a future pass could search for it directly, the way section 22.12 eventually did for Schwarz's 1869 paper. What remains open is exactly what remained open before: neither this source nor section 5.8's reproduces the actual harmonic-extension argument step by step, so this volume still has the technique's *name* and general shape rather than a checked derivation, and theorem 5.10's own geodesic-disk theorem (as opposed to its Chebyshev exposition) is untouched by either source, exactly as section 5.8's closing paragraph already states.

What section 5.8 does not supply is any part of Milnor's own geodesic-disk theorem (theorem 5.10 as stated in section 5.6 above) beyond what section 5.6's original two secondary sources already established: Papadopoulos's detailed treatment is of the Chebyshev exposition specifically, and his account of Milnor's own azimuthal-equidistant result is a single summary sentence, not a restated theorem. Items (a) and (c) of the surviving todobox above are accordingly unchanged as they apply to theorem 5.10 itself; only the Chebyshev-theorem portion of item (a) — whether the stated hypotheses and uniqueness clause are exactly as Milnor gave them — is now settled, and only for that portion. A direct reading of Milnor's 1969 *American Mathematical Monthly* text remains the only route to closing what is left, and that text itself was not obtained in this pass: it sits behind a Taylor & Francis subscription wall (DOI `10.1080/00029890.1969.12000424`), a different and more conventional access barrier than the scanned-primary-source problem this volume has met repeatedly elsewhere (section 16.3, appendix A.5).

## 5.9 How the two results of this chapter relate

It is worth being explicit that theorem 5.2 is not a special case, a corollary, or an independent proof of Milnor's theorem; it is a different theorem, about a different functional (a boundary

flux of $\ln \mu$, rather than the global scale-ratio spread $\delta(f)$ of definition 5.9), proved by a different method (Gauss–Bonnet applied to the curvature identity of lemma 5.1, derived directly in this chapter from eq. (2.4) rather than quoted, after section 5.2's correction of the version originally quoted from chapter 3), established here with full rigor for the restricted case of conformal projections only, whereas theorem 5.10 holds for arbitrary homeomorphisms. The two results are nonetheless visibly compatible rather than merely coincidentally similar: both identify the disk-shaped region (a geodesic cap, or the topological disk $U$ of theorem 5.2) as the natural test case, both find that distortion vanishes to leading order as the region shrinks, and corollary 5.4's pointwise compression result, once the boundary is true-to-scale, is the same qualitative phenomenon that theorem 5.10 makes sharp for the worst-case scale ratio specifically. Whether theorem 5.2 can be derived as a direct corollary of theorem 5.10 restricted to conformal maps, is genuinely independent of it, or is subsumed by a version of it not stated here, is not yet known and is not claimed either way; resolving this is one of the items left for the primary-source check recorded in section 5.6's todobox, alongside the still-unverified optimality half of theorem 5.10 itself. What this chapter has established without qualification is theorem 5.2 and, building directly on it, lemma 5.3 and corollary 5.4 of section 5.3; what it has established with a specific, sourced, but not primary-verified statement is theorem 5.10; and what remains entirely open, for both, is a unified single theorem covering global distortion for the general (not necessarily conformal) case with a proof reproduced in this volume rather than cited from either source. It is worth flagging explicitly that corollary 5.4's pointwise statement and theorem 5.10's $\sigma_1$ are not known to coincide even for boundary-true-to-scale conformal projections of a geodesic disk: $\sigma_1$ is an infimum over *finite* point-pair scales in definition 5.9's sense, whereas corollary 5.4 bounds the *pointwise infinitesimal* scale $\mu$, and section 5.6's own remark that these two notions "typically... coincide, though this requires an argument not given here" is exactly as unproven now as it was before this section; corollary 5.4 narrows what such an argument would need to show but does not supply it. Section 5.10 below closes this specific item.

## 5.10 Closing the gap: $\sigma_1, \sigma_2$ coincide with the pointwise extremes, on a geodesically convex region

The remark left standing at the end of section 5.9 is, unlike the items in section 5.6's todobox, not a primary-source question at all: it is a self-contained claim about how a finite-distance quantity built from geodesic and Euclidean distances relates to a pointwise-infinitesimal one, and it can be settled entirely from material already in this volume — the Tissot indicatrix machinery of chapter 3, giving at each point of a smooth map the two principal scales $a(P) \geq b(P) > 0$ bounding the local expansion factor in every direction — without consulting Milnor's paper or any other primary source.

**Definition 5.12** (Geodesic convexity)**.** A region $U$ of the sphere is *geodesically convex* if, for every pair of distinct points $x, y \in U$, some minimizing geodesic from $x$ to $y$ (a great-circle arc of length $d_S(x, y)$) lies entirely within $U$.

**Key Result**

*Every geodesic disk $D_\alpha$ with $\alpha \le \pi/2$ is geodesically convex in this sense: two points of $D_\alpha$ at angular distance $\le \pi$ apart are joined by a unique minimizing great-circle arc whenever they are not antipodal, and elementary spherical trigonometry (the triangle inequality for geodesic distance, applied to the centre $O$ and a point $Q$ on the joining arc) gives $d_S(O,Q) \le \max(d_S(O,x), d_S(O,y)) \le \alpha$ for every $Q$ on the minimizing arc between $x, y \in D_\alpha$, so the arc does not leave $D_\alpha$; no such elementary argument is available for $\alpha > \pi/2$, where two points can lie on opposite sides of a great circle through $O$ close enough to being antipodal that their minimizing arc bulges outside $D_\alpha$, so the hypothesis below is stated for general geodesically convex $U$ rather than asserted for every disk in this chapter.*

**Lemma 5.13** (Length distortion along a path is controlled by the pointwise extremes)**.** *Let $f$ be a $C^1$ diffeomorphism of a region $U$ of the sphere onto a region of the Euclidean plane, with principal scales $a(P) \ge b(P) > 0$ at each $P \in U$ in the sense of chapter 3 (so that the local linearization of $f$ at $P$ maps a unit tangent vector making angle $\eta$ with the major axis of the indicatrix to a vector of Euclidean length between $b(P)$ and $a(P)$, and every intermediate value is attained as $\eta$ varies). Let $\gamma : [0, L] \to U$ be any $C^1$ path parametrized by sphere-arclength ($L = \mathrm{length}_S(\gamma)$), and let $f \circ \gamma$ be its image. Then*

$$\inf_{P \in \gamma} b(P) \cdot L \;\le\; \mathrm{length}_E(f \circ \gamma) \;\le\; \sup_{P \in \gamma} a(P) \cdot L.$$

*Proof.* By the chain rule, $(f\circ\gamma)'(s) = df_{\gamma(s)}(\gamma'(s))$, and since $\gamma$ is unit-speed, $|\gamma'(s)|_S = 1$, so $|(f\circ\gamma)'(s)|_E$ is exactly the local scale of $f$ at $\gamma(s)$ in the direction $\gamma'(s)$, which by definition of $a, b$ (chapter 3) lies in $[b(\gamma(s)), a(\gamma(s))]$. Integrating this pointwise bound over $s \in [0, L]$, $\mathrm{length}_E(f \circ \gamma) = \int_0^L |(f \circ \gamma)'(s)|_E \, ds$ is sandwiched between $\int_0^L b(\gamma(s))\, ds \ge \inf_\gamma b \cdot L$ and $\int_0^L a(\gamma(s))\, ds \le \sup_\gamma a \cdot L$. □

**Theorem 5.14** (Milnor's $\sigma_1, \sigma_2$ equal the pointwise extremes of the principal scales, on a geodesically convex region)**.** *Let $U$ be a geodesically convex region of the sphere and $f : U \to \mathbb{R}^2$ a $C^1$ diffeomorphism onto its image, with principal scales $a(P) \ge b(P) > 0$ as in lemma 5.13, and suppose $f^{-1}$ extends to a $C^1$ diffeomorphism of the image region as well (automatic if $f$ is, e.g., conformal with $\mu > 0$ everywhere, the case of interest in this chapter). Then, with $\sigma_1, \sigma_2$ as in definition 5.9,*

$$\sigma_1 \;=\; \inf_U b(P), \qquad \sigma_2 \;=\; \sup_U a(P).$$

*Proof. Upper bound on $\sigma_2$.* For distinct $x, y \in U$, geodesic convexity provides a minimizing geodesic $\gamma$ from $x$ to $y$ lying in $U$, of sphere-length $d_S(x, y)$. Since a straight line in the plane is the shortest path between its endpoints, $d_E(f(x), f(y)) \le \mathrm{length}_E(f \circ \gamma)$, and lemma 5.13 bounds the right side by $\sup_U a \cdot d_S(x, y)$. Hence $d_E(f(x), f(y))/d_S(x, y) \le \sup_U a$ for every pair, giving $\sigma_2 \le \sup_U a(P)$.

*Lower bound on $\sigma_1$.* Apply the same argument to $f^{-1}$, whose own principal scales at $f(P)$ are exactly $1/a(P), 1/b(P)$ (the local linearization of $f^{-1}$ at $f(P)$ is the inverse linear map of $df_P$, whose singular values are the reciprocals of $df_P$'s). The image region is geodesically convex in the trivial (Euclidean) sense that every two points are joined by the straight segment between them, which lies in any convex set and in particular need not be checked further here since Euclidean straight-line paths always exist and are always the shortest. Applying lemma 5.13 to $f^{-1}$ along the straight segment $\gamma_E$ from $f(x)$ to $f(y)$ (of Euclidean length $d_E(f(x), f(y))$) gives $\mathrm{length}_S(f^{-1} \circ \gamma_E) \le \sup(1/b) \cdot d_E(f(x), f(y)) = d_E(f(x), f(y))/\inf_U b$. Since $f^{-1} \circ \gamma_E$ is some path from $x$ to $y$ on the sphere, its sphere-length is at least the geodesic distance: $d_S(x, y) \le \mathrm{length}_S(f^{-1}\circ\gamma_E) \le d_E(f(x), f(y))/\inf_U b$. Rearranging, $d_E(f(x), f(y))/d_S(x, y) \ge \inf_U b$ for every pair, giving $\sigma_1 \ge \inf_U b(P)$.

*Matching lower bound on $\sigma_2$ and upper bound on $\sigma_1$.* Fix $\varepsilon > 0$ and a point $P_0 \in U$ with $a(P_0) > \sup_U a - \varepsilon$ (or, if the supremum is attained, take $P_0$ to be a maximizer). Let $\eta_0$ be the direction of the major axis of the indicatrix at $P_0$, and let $y_t$ be the point at geodesic distance $t$ from $P_0$ along the geodesic in direction $\eta_0$, for small $t > 0$ (contained in $U$ for $t$ small since $U$ is open, or, at worst, up to the boundary). As $t \to 0^+$, $d_E(f(P_0), f(y_t))/d_S(P_0, y_t) \to a(P_0)$ directly from the definition of $a$ as the local scale in the direction of maximal expansion (chapter 3: this is exactly the defining limit of the principal scale, the same limit used to define $\mu$ itself in lemma 5.1's proof for the conformal case). Hence $\sigma_2 = \sup_{x \neq y}(\cdot) \geq \limsup_{t \to 0^+}(\cdot) = a(P_0) > \sup_U a - \varepsilon$; letting $\varepsilon \to 0$, $\sigma_2 \geq \sup_U a(P)$. Combined with the upper bound above, $\sigma_2 = \sup_U a(P)$ exactly. The symmetric argument at a near-minimizer of $b$, using the minor-axis direction, gives $\sigma_1 \leq \inf_U b(P) + \varepsilon$ for every $\varepsilon > 0$, hence $\sigma_1 \leq \inf_U b(P)$; combined with the lower bound above, $\sigma_1 = \inf_U b(P)$ exactly. □

**Key Result**

*Theorem 5.14 closes the item flagged at the end of section 5.9, in the case that matters for this chapter's own results: every conformal map has $a(P) = b(P) = \mu(P)$ at each point (chapter 4), so for a conformal $f$ on a geodesically convex $U$, theorem 5.14 gives directly $\sigma_1 = \inf_U \mu$ and $\sigma_2 = \sup_U \mu$, with no residual gap and no "typically" qualifier. Applied to theorem 5.6's exact radial profile on $D_\alpha$ for $\alpha \leq \pi/2$ (within the geodesically convex range definition 5.12 needs), this identifies Milnor's own finite-distance quantities exactly: $\sigma_1 = \mu(0) = \cos^2(\alpha/2)$ and $\sigma_2 = \mu(\alpha) = 1$, so that definition 5.9's distortion for* this specific conformal, boundary-normalized construction *is $\delta(f) = \log(1/\cos^2(\alpha/2)) = -2\ln\cos(\alpha/2)$ — exactly $\delta_{\text{conf}}(\alpha)$ as already computed after theorem 5.6, now known to be the literal value of Milnor's own functional for this map rather than an analogue of it justified only by informal appeal to "typically coincide." This does not touch theorem 5.10 itself, whose content is a claim about the infimum of $\delta(f)$ over* all *homeomorphisms, conformal or not, and whose optimality half remains exactly as unverified against Milnor's primary text as section 5.8's todobox already states; what changes is that the gap identified after theorem 5.6 between $\delta_{\text{conf}}(\alpha)$ and $\delta_0(\alpha)$ is now known to be a gap between two literally commensurable quantities — Milnor's own $\delta$, evaluated at a specific conformal map versus minimized over all maps — rather than between $\delta_0(\alpha)$ and a merely analogous, unidentified substitute.*

# Chapter 6

# The Hairy Ball Theorem: A Topological, Not Metric, Obstruction

## 6.1 A different kind of impossibility

Every impossibility result so far in this volume has had the same source: the Gaussian curvature of the sphere is non-zero (chapter 2), and this metric fact propagates, via Tissot's construction (chapter 3), into the pointwise (chapter 4) and global (chapter 5) impossibility of certain distortion-free projections. This chapter's obstruction is of a different character entirely: it does not depend on the sphere's curvature at all, and would hold equally for a badly dented, non-round sphere of any metric whatsoever, because it is a fact about the sphere's *topology* — specifically, about the impossibility of a certain kind of continuous vector field — rather than about lengths, angles, or areas. The connection to cartography is nonetheless direct and, for the purposes of this series, urgent: the retroazimuthal projections at the centre of *The Retroazimuthal Problem* are constructed, at root, from a field of preferred directions on the sphere (the direction, at each point, of the great-circle bearing toward a fixed centre), and this chapter supplies the topological reason that construction cannot be extended, unmodified, to the whole sphere at once.

## 6.2 Vector fields on a surface

**Definition 6.1** (Tangent vector field)**.** A *tangent vector field* on a surface $S \subset \mathbb{R}^3$ is a continuous assignment $P \mapsto v(P) \in T_P S$ of a vector in the tangent plane at $P$ to each point $P \in S$. The field is *nowhere zero* if $v(P) \neq 0$ for every $P$.

On the plane, or on any open disk, nowhere-zero tangent vector fields obviously exist in abundance: the constant field $v \equiv (1, 0)$ is one. The question this chapter answers is whether the same is true on the whole sphere, and the answer — no — is the celebrated *hairy ball theorem*, informally: one cannot comb a hairy ball flat without a cowlick.

**Theorem 6.2** (Hairy ball theorem, Poincaré 1885 / Brouwer 1912)**.** *Every continuous tangent vector field on the sphere $S^2$ has at least one zero. Equivalently, there is no continuous, nowhere-zero tangent vector field on $S^2$.*

**Epistemic Notice**

The attribution above follows the standard account in the topology literature: Poincaré's 1885 memoir on curves defined by differential equations contains, in essence, the index-sum argument for vector fields on surfaces from which the theorem for $S^2$ follows as a special case, and Brouwer's 1912 paper gave the modern, fully general proof via degree theory for arbitrary even-dimensional spheres. Attaching precise priority (whether Poincaré's 1885 statement already covers the $S^2$ case explicitly, or whether it is only an immediate corollary later drawn out) is not something this draft has verified against either primary text and is flagged for the same reason as the analogous historical notes in chapters 2 and 3.

### 6.2.1 Proof via degree theory

The proof given here is the standard one via the topological degree of a map, chosen because the same technique — tracking how many times a map wraps around, and observing that this integer cannot change under continuous deformation — is exactly the technique needed, in section 6.3 below, to see *why* the retroazimuthal direction field's obstruction is not a defect of one particular construction but a necessary feature of any construction of the same general kind.

*Proof sketch.* Suppose, for contradiction, that $v$ is a continuous nowhere-zero tangent vector field on $S^2$. Normalize it to a unit field, $u(P) = v(P)/|v(P)|$, still continuous and nowhere zero.

*Step 1: build a homotopy from the identity to the antipodal map.* For each $P \in S^2$ and $t \in [0, \pi]$, rotate $P$ within the great circle determined by $P$ and $u(P)$ (the tangent direction gives a well-defined instantaneous rotation axis at $P$, namely the axis perpendicular to both $P$ and $u(P)$) through angle $t$; call the result $H(P, t)$. This is continuous in $(P, t)$ because $u$ is continuous and nowhere zero (so the rotation axis is always well defined — this is exactly where nowhere-vanishing is used), $H(P, 0) = P$, and $H(P, \pi) = -P$, the antipodal point, since rotating by $\pi$ about an axis perpendicular to $P$ sends $P$ to $-P$ regardless of which such axis is used.

*Step 2: compare degrees.* $H(\cdot, 0)$ is the identity map $S^2 \to S^2$, of topological degree $+1$. $H(\cdot, \pi)$ is the antipodal map $P \mapsto -P$, of topological degree $(-1)^3 = -1$ (the antipodal map on $S^n$ has degree $(-1)^{n+1}$; for $n = 2$, degree $-1$). But $H$ is a continuous homotopy between these two maps, and topological degree is a homotopy invariant (this is the key black-box fact from degree theory used here, itself a standard consequence of the fact that degree is computed by an integral or intersection-number formula invariant under continuous deformation of the map). Hence $+1 = -1$, a contradiction. □

**Key Result**

*The proof shows something slightly sharper than the bare existence statement: it shows the obstruction is exactly the mismatch between the degree of the identity map ($+1$) and the degree of the antipodal map ($-1$) on $S^2$, a mismatch that occurs precisely because 2 is even (on odd-dimensional spheres $S^{2k+1}$, the antipodal map has degree $(-1)^{2k+2} = +1$, matching the identity, and indeed nowhere-zero tangent fields* do *exist on every odd-dimensional sphere — e.g. $(x, y, z, w) \mapsto (-y, x, -w, z)$ on $S^3$). This is why the obstruction is genuinely topological rather than a low-dimensional accident: it is controlled by the Euler characteristic $\chi(S^2) = 2 \neq 0$, via the Poincaré–Hopf theorem, of which theorem 6.2 is the special case for a nowhere-zero field (Poincaré–Hopf: the sum of the indices of the zeros of any vector field on a closed surface equals $\chi$; if $\chi \neq 0$ there must be at least one zero, of nonzero index).*

## 6.3 Application: why a global retroazimuthal direction field cannot exist

A *retroazimuthal* projection, in the sense developed at length in *The Retroazimuthal Problem* and only briefly recalled here, is one with a distinguished centre point $C$ (in that volume's central application, the location of the Kaaba, for the specific cartographic problem of the qibla direction) such that the direction from any point $P$ toward $C$, in the map's own planar coordinates, equals the true compass bearing of the great-circle route from $P$ to $C$ on the sphere. Chapter 12 of that companion volume ("A Common Vector-Field Structure Behind Craig, Hammer, and Littrow") does formulate the retroazimuthal condition as an alignment between a map's planar bearing and a true-bearing vector field $V(\varphi, \lambda)$, and derives a scalar factor $k$ making the three named historical constructions particular cuts of one family; that vector-field formulation is genuine background for the construction of $\beta(P)$ below. It does *not*, however, contain any singularity, index, or topological argument: the following construction, and the index computation carried out at $C$ and $-C$ below, are original to this chapter, not recalled from the companion volume.

**Proposition 6.3** (No globally consistent retroazimuthal direction field)**.** *Let $C \in S^2$ be fixed. The bearing field $\beta(P)$, defined for $P \in S^2 \setminus \{C, -C\}$ as the initial unit tangent direction of the great-circle path from $P$ to $C$, cannot be extended to a continuous, nowhere-zero tangent vector field on the whole sphere $S^2$ (including at $C$ and $-C$).*

*Proof.* At $P = C$ itself, every direction is simultaneously "toward $C$" in the degenerate sense of distance zero, and correspondingly $\beta(P)$ has no continuous limit as $P \to C$ from different directions: approaching $C$ along different meridians of the local geodesic-normal coordinate system gives limiting bearings pointing in every possible tangent direction at $C$, so no single continuous extension of $\beta$ to $P = C$ exists; the field genuinely has a singularity there, of the same qualitative type as the "centre" singularity of the standard planar radial field $(x, y)/\sqrt{x^2 + y^2}$ at the origin, which has index $+1$. At $P = -C$, the antipode, every great-circle path to $C$ is equally short (all meridians through $-C$ reach $C$ after the same distance $\pi R$), so again every direction is simultaneously a valid limiting bearing, and again no continuous extension exists; this second singularity is of index $-1$. (No index computation of this kind appears in *The Retroazimuthal Problem*; the identification of these two singularities and their indices is original to this chapter, not recalled from the companion volume — see the corrected todobox below.) Since these two necessary singularities have indices summing to 0, this proposition's failure of extension is not yet, by itself, a direct instance of theorem 6.2 (a field with a $+1$ and a $-1$ singularity is topologically consistent with the Poincaré–Hopf sum $\chi(S^2) = 2$ only if there are further singularities elsewhere accounting for the remaining index $+2$, or if the two singularities identified are not the whole story); the substantive claim, proved directly rather than merely asserted, is that $\beta$ is smooth and nowhere zero away from $\{C, -C\}$ (immediate: the great-circle bearing is a well-defined, smoothly-varying unit tangent vector at every point other than the two singular points, since geodesics on the sphere vary smoothly with endpoints away from the cut locus) so that *if* a continuous nowhere-zero extension to all of $S^2$ existed, it would contradict theorem 6.2 directly. Hence no such extension exists, and the singularities at $C, -C$ identified above are exactly the two zeros (or points of non-existence, made into zeros by any continuous completion) that theorem 6.2 guarantees must occur somewhere. □

**Key Result**

*Proposition 6.3 is the precise topological statement underlying, but not already stated in,* The Retroazimuthal Problem*'s case-by-case treatment of Craig's, Hammer's, and Littrow's constructions, each of which is indeed restricted to a proper subset of the sphere, though not by a single uniform pattern: Craig's construction and Littrow's conformal construction are each defined away from small neighbourhoods of $C$ and of the antipode $-C$ (and, for Littrow, of the poles as well); Hammer's construction is instead defined* separately *on the front hemisphere ($|\lambda_{rel}| \leq 90°$) and the back hemisphere, with the two combined images overlapping non-injectively unless the back hemisphere is corrected by a* 180° *rotation; and the oblique conformal construction generalizing Littrow, treated in that volume's Chapter 13, is shown by direct complex-analytic argument (not by an index computation) to have exact domain $S^2 \setminus \{C, \bar{C}\}$. None of these three, individually or together, amounts to the topological argument given above — proposition 6.3 and theorem 6.2 are what* explain*, formula-independently, why every one of these constructions was forced to exclude at least $C$ and $-C$ no matter how cleverly designed, but that explanation is not itself present in the companion volume and is original to this chapter.*

**Cross-reference verified against the primary text, 1 August 2026.** The cross-reference to *The Retroazimuthal Problem*, Chapter 12, flagged by the previous todobox, has been checked line-by-line against that chapter's actual text (`ch09b_struttura_unificata.tex` in the companion project, the twelfth `\chapter` in reading order). The check found the original claim to be **incorrect on two points** and has corrected the text above accordingly: (a) Chapter 12 contains no index computation and no singularity analysis of any kind for the field $\beta$ at $C$ and $-C$; a full-text search of the companion volume's source for *topological*, *degree theory*, *winding number*, *Poincaré*, and *hairy ball* returns zero matches anywhere in the volume, confirming the topological argument of this chapter is genuinely new rather than recalled. (b) The domain-of-definition claim attributed to Chapter 12 was also wrong in two ways: first, the relevant material is not in Chapter 12 at all but scattered across Chapters 9, 8, and 6 (Craig, Hammer, Littrow individually) and Chapter 13 (the oblique conformal generalization, where the domain $S^2 \setminus \{C, \bar{C}\}$ is established by Cauchy–Riemann and complex-analytic argument, not an index computation); second, the "hemisphere or sphere-minus-antipode" dichotomy was too coarse — Hammer's construction in particular requires two separately-defined hemispheres with an explicit non-injectivity correction, not a single excluded antipode. Chapter 12 itself (the vector-alignment criterion and the scalar field $k$ unifying Craig, Hammer, and Littrow) is legitimate and relevant background, now cited correctly above, but is unrelated to the singularity/index content. Item (c) of the original todobox is resolved: proposition 6.3 is not stated, in this or an equivalent topological form, anywhere in the companion volume, so it is presented here as new.

## 6.4 What remains for item (iv)

Proposition 6.3 rules out a globally defined retroazimuthal *direction field*, which is the topological obstruction referenced as item (iii) in chapter 1's list. Item (iv) — that no projection is simultaneously conformal and retroazimuthal with an arbitrary prescribed centre, beyond the cases already classified in *The Retroazimuthal Problem* — is a separate, sharper claim: even restricted to a single hemisphere where proposition 6.3's obstruction does not apply, requiring *both* retroazimuthality and conformality simultaneously is a further, metric (not merely topological) constraint, of the same general flavour as chapter 4's conformal/equal-area result but for a different pair of properties, and it is not derived from theorem 6.2 at all. That argument — planned as chapter 7 in chapter 1's outline — is not attempted in this draft; it requires the Cauchy–Riemann-type analysis already flagged, in the introductory chapter, as a live theme carried over from the companion volume, and is left for a future chapter of this monograph.

# Chapter 7

# Item (iv), Revisited: Conformality and Retroazimuthality Are Compatible, and the Only Obstruction Is Topological

**Epistemic Notice**

This chapter was planned, in chapter 1, as a fourth impossibility result: that no map is simultaneously conformal and retroazimuthal with an arbitrarily prescribed centre, "beyond the cases already classified." Actually carrying out the verification this volume's discipline requires — checking the claim against the companion volume's own results, not against a recollection of them — shows that this planned result is **false as stated**. *The Retroazimuthal Problem*, Chapters 12 and 13, already establish, by independent symbolic and numerical verification, that a conformal retroazimuthal projection *exists* for every centre, not only the equatorial one Littrow's own formula was built for. This chapter reports that correction rather than the originally planned proof, consistent with the no-invented-data discipline stated in chapter 1: a chapter is not written to match a plan drawn up before the relevant checking was done.

## 7.1 What item (iv) originally claimed

The four-item programme of chapter 1 listed, as its fourth and sharpest impossibility claim, that no map from a neighbourhood of a point on the sphere to the plane is simultaneously conformal and retroazimuthal with an arbitrarily prescribed centre. The intuition behind the claim, at the time it was planned, was that Littrow's projection — the sphere's one classically known conformal retroazimuthal construction — works only for a centre on the equator of its own coordinate system, and that generalising it to an arbitrary centre should fail for a structural reason connected to the topological obstruction of chapter 6.

## 7.2 What the companion volume actually establishes

*The Retroazimuthal Problem* already carried out, independently and before the present volume existed, exactly the analysis needed to settle this question, in two stages.

### 7.2.1 Chapter 13: the retroazimuthal condition as a Cauchy–Riemann system

Fixing a centre $C = (\varphi_C, \lambda_C)$ and writing a candidate map in polar form around $f(C)$, with $r = e^s$ the unknown radial scale and $\theta(P)$ the (known, elementary) true bearing from $P$ to $C$, conformality of $f$ is shown there to be equivalent to the linear system

$$s_\lambda = -\theta_\psi, \qquad s_\psi = \theta_\lambda,$$

in the isothermal coordinates $(\lambda, \psi)$, $\psi = \operatorname{artanh}(\sin\varphi)$. Local solvability of this system for $s$ requires $\theta$ to be *harmonic*, $\theta_{\lambda\lambda} + \theta_{\psi\psi} = 0$. That volume checks this symbolically in SymPy for symbolic $\varphi_C$, not only at $\varphi_C = 0$, and reports it as identically zero for every centre tested.

This specific claim has been re-verified independently for this chapter, not merely taken on trust from the companion volume: with $V_x = -\sin\lambda\cos\varphi_C$, $V_y = \cos\varphi\sin\varphi_C - \sin\varphi\cos\varphi_C\cos\lambda$ and $\theta = \operatorname{atan2}(V_x, V_y)$, SymPy's `simplify` reduces $\theta_{\lambda\lambda} + \cos\varphi\,\partial_\varphi(\cos\varphi\,\theta_\varphi)$ (the isothermal Laplacian rewritten in $\varphi$) to the literal expression `0`, for fully symbolic $\varphi_C$, and separately at three numerical spot-checks ($\varphi_C = 0$, $\varphi_C = \pi/6$, and $\varphi_C \approx 21.389°$, Mecca's latitude), all confirming exact zero. This matches the companion volume's own reported result and is not a new finding, but it is now a finding this volume has checked for itself rather than only recalled.

**Key Result**

*Harmonicity of $\theta$ — the local, necessary condition for a conformal retroazimuthal construction to exist near a given centre — holds* for every centre, *not only the equator. There is no local, differential obstruction that could explain why an arbitrary centre should behave differently from Littrow's equatorial one. Whatever distinguishes Littrow's centre, if anything genuinely does, cannot be visible at this level.*

### 7.2.2 Chapters 12–13: the local obstruction does not survive globally, and a closed form exists

The companion volume goes further than the local (harmonicity) check. Its Chapter 13 computes the monodromy period of the candidate $s$ around a loop generating $\pi_1(\mathbb{S}^2 \setminus \{C, \bar{C}\}) = \mathbb{Z}$ and finds it numerically zero, to floating-point precision, for eight tested centres between the equator and 80° latitude, including Mecca's. Its Chapter 12 then closes the loop analytically rather than only numerically: writing the retroazimuthal condition as a vector-alignment $(-x, -y) = k(\varphi, \lambda)\,(V_x, V_y)$ (that chapter's general vector-alignment criterion), it shows *both* that the conformal (Littrow-type) member of this family necessarily has $k = \sec\varphi$, for every centre $\varphi_C$, and — the converse, established by direct symbolic substitution rather than by any further numerical sampling — that $k = \sec\varphi$ substituted back into the vector-alignment equation reproduces the volume's closed-form conformal retroazimuthal construction *exactly*, for arbitrary symbolic $\varphi_C$, with no appeal to complex analysis in that final step. Between the two chapters, existence for an arbitrary centre is established twice, by independent routes (a Cauchy–Riemann/monodromy argument in Chapter 13, and a direct algebraic substitution in Chapter 12), that agree.

**Epistemic Notice**

This existence result is not re-derived in this chapter; it is recalled, with the specific propositions and verification method identified, from a companion volume of the same author that this volume treats as an already-verified external source, in the same way chapter 6 recalls the vector-alignment criterion of that volume's Chapter 12 for a different purpose. Only the harmonicity step above has been independently re-run for this chapter specifically.

## 7.3 The domain restriction is forced by topology, not by a separate metric obstruction

The construction recalled in section 7.2 is defined on $\mathbb{S}^2 \setminus \{C, \bar{C}\}$, not on the whole sphere: $\theta(P)$ itself is undefined at $P = C$ (every direction is simultaneously "toward $C$") and at $P = \bar{C}$ (every direction is equally a shortest path to $C$), exactly the two points identified, by an unrelated topological argument, as the necessary singularities of any retroazimuthal bearing field in proposition 6.3 of chapter 6.

**Key Result**

*This is the genuine content that survives from the original plan for item (iv), and it is what this chapter's title records: the domain restriction to $\mathbb{S}^2 \setminus \{C, \bar{C}\}$ is not a separate, additional metric obstruction sitting alongside conformality, waiting to be discovered by some future refinement of the Cauchy–Riemann analysis above. It is exactly, and only, chapter 6's topological obstruction, already fully proved there without reference to conformality at all. A conformal retroazimuthal construction is, by section 7.2, achievable everywhere the underlying bearing field $\beta$ is itself defined and smooth — and nowhere else, because $\beta$ genuinely cannot be extended past $C, \bar{C}$, conformally or otherwise. The metric (Cauchy–Riemann) and topological (hairy ball) pictures, developed independently in two different chapters of two different volumes, are found here to coincide exactly on where the obstruction lives.*

## 7.4 What remains open

**Todo — verify primary source**

Two questions were left unresolved by the recalled results. First, uniqueness *outside* the vector-alignment family used in Chapter 12 of the companion volume: section 7.2 shows that $k = \sec\varphi$ is necessary and sufficient for conformality *within* that specific one-parameter construction, but did not rule out some structurally different conformal retroazimuthal map, not built from a scalar multiple of the bearing vector $V$ at all, existing on the same punctured domain with different qualitative behaviour. Section 7.5 below closes this question. Second, the secondary-literature uniqueness claim attributed to Maurer (1919) — that Littrow's construction is the *only* conformal retroazimuthal projection, full stop — was already flagged in the companion volume as unverified against Maurer's original text; the results recalled above make that claim harder to sustain as stated (since a conformal retroazimuthal construction is now exhibited for every centre, not only the equatorial one), but resolving the apparent tension with Maurer's claim requires either obtaining Maurer's paper or a citation tracing exactly what he proved, neither of which this chapter attempts.

## 7.5 Uniqueness beyond the vector-alignment family, closed

Section 7.2's todobox left open whether some conformal retroazimuthal map, not built as a scalar multiple of the bearing vector field $V$ the way the companion volume's Chapter 12 family is, might exist on the same punctured domain with genuinely different qualitative behaviour. This section closes that question: no such map exists, because the retroazimuthal condition itself — not the vector-alignment ansatz, which turns out not to be a restriction at all — already forces every candidate into the one-parameter family already analyzed.

**Proposition 7.1** (The retroazimuthal condition forces the polar form; it is not chosen)**.** *Let $f$ be any map, conformal or not, defined on a domain $U \subseteq S^2 \setminus \{C, \bar{C}\}$ containing an open set, such that $f(C)$ extends to a well-defined finite point in the plane and $f$ is retroazimuthal with centre $C$ in the sense of chapter 6's definition: for every $P \in U$, the direction of the*

*straight segment from $f(P)$ to $f(C)$ in the plane equals the true compass bearing $\theta(P)$ of the great-circle route from $P$ to $C$. Then, in polar coordinates $(r, \alpha)$ centred at $f(C)$, $f(P)$ has angular coordinate $\alpha(P) = \theta(P)$ (up to a single additive constant $\alpha_0$ fixed once for the whole map, the choice of which planar direction represents "north"), for every $P \in U$.*

*Proof.* This is immediate from the definition recalled in chapter 6: the retroazimuthal condition is stated directly as a condition on the *direction* of $f(P)$ as seen from $f(C)$, and the angular coordinate of a point in polar coordinates centred at $f(C)$ is, by definition of polar coordinates, exactly that direction (measured against whatever fixed reference ray the coordinate system's $\alpha = 0$ represents). Choosing that reference ray to coincide with the map's own "north" direction gives $\alpha_0 = 0$; any other choice of reference ray shifts every $\alpha(P)$ by the same constant $\alpha_0$, with no other effect. There is no step in this argument at which a choice is made: the hypothesis *is* the polar-angle condition, restated in polar coordinates rather than as a condition on a line's direction. □

**Proposition 7.2** (A holomorphic function is determined by its real part, up to an additive imaginary constant)**.** *Let $g_1, g_2$ be holomorphic on a connected open set $\Omega \subseteq \mathbb{C}$ with $\operatorname{Re} g_1 \equiv \operatorname{Re} g_2$ on $\Omega$. Then $g_1 - g_2$ is a purely imaginary constant on $\Omega$.*

*Proof.* Write $h := g_1 - g_2 = u + iv$, holomorphic with $u \equiv 0$ on $\Omega$. The Cauchy–Riemann equations $u_x = v_y$, $u_y = -v_x$ give $v_y \equiv 0$ and $v_x \equiv 0$ on $\Omega$, so $v$ is locally constant; since $\Omega$ is connected, $v$ is a single constant $c \in \mathbb{R}$ throughout. Hence $h \equiv ic$. □

**Theorem 7.3** (Uniqueness of the conformal retroazimuthal map, up to the similarity group)**.** *Fix a centre $C$ and let $w = \lambda + i\psi$ be the isothermal complex coordinate of section 7.2. Suppose $f_1, f_2$ are both conformal retroazimuthal maps with centre $C$, orientation-preserving, defined and single-valued on all of $S^2 \setminus \{C, \bar{C}\}$ (equivalently, on the corresponding cylinder in the $w$-coordinate). Then $f_2 = k \cdot f_1$ for some real constant $k > 0$ (a uniform rescaling): the two differ only by choice of overall map scale.*

*Proof.* By proposition 7.1, both $f_1, f_2$ have angular coordinate $\theta(P)$ around their respective centres $f_1(C), f_2(C)$ (taking the same reference-north convention for both, so $\alpha_0 = 0$ in each case; this is a labelling choice, not a substantive one, and is fixed once for the comparison). Write $f_i(P)$ in the polar form $r_i(P) = e^{s_i(P)}$, angle $\theta(P)$, for $i = 1, 2$; this is exactly section 7.2's ansatz, now shown to be forced rather than assumed. Conformality of $f_i$ is, by that section's recalled result, equivalent to the Cauchy–Riemann system for $g_i(w) := \theta(w) + is_i(w)$, i.e. $g_i$ is holomorphic in $w$ on the domain in question. Both $g_1, g_2$ are holomorphic with $\operatorname{Re} g_1 = \operatorname{Re} g_2 = \theta$ (the same true-bearing function, since both maps share the same centre $C$ and the underlying sphere geometry $\theta$ depends only on $C$, not on $f$), and the domain (the $w$-image of $S^2 \setminus \{C, \bar{C}\}$, an infinite cylinder, connected) is connected. By proposition 7.2, $g_1 - g_2 \equiv ic$ for some real constant $c$, i.e. $s_1 - s_2 \equiv c$ identically. Hence $r_2(P)/r_1(P) = e^{s_2(P) - s_1(P)} = e^{-c} =: k$ is the same positive constant at every point, and since the angular coordinates already agree exactly, $f_2(P) = k \cdot f_1(P)$ (as vectors from the shared centre point) for every $P$. □

**Epistemic Notice**

Three assumptions in theorem 7.3's statement are worth making explicit rather than leaving implicit in the proof, in keeping with this volume's practice of stating exactly what a uniqueness result does and does not cover.

*Orientation.* An orientation-*reversing* conformal map (built from $\bar{w}$ rather than $w$) is not excluded by the retroazimuthal condition alone in the abstract — proposition 7.1's argument does not use orientation at all — but is excluded by requiring $f$ to reproduce the *actual* compass bearing $\theta(P)$ rather than its mirror image $-\theta(P)$ under a reflected map convention; this is the standard cartographic convention (a printed map is not usually read in a mirror), stated here as a hypothesis rather than smuggled in silently. Dropping it would add exactly one further discrete choice (reflection) to the similarity group in theorem 7.3's conclusion, not a structurally new family.

*Single-valuedness.* The proof does not construct $s_i$ from $\theta$ by integration — which would require checking the domain's nontrivial fundamental group does not obstruct single-valuedness, exactly section 7.2's monodromy computation — but instead compares two *already single-valued* candidates $g_1, g_2$, for which proposition 7.2's argument needs only connectedness, not simple connectivity. This is a genuine simplification available specifically because the question posed is uniqueness among valid maps, not existence: existence needs the harder monodromy fact (already established); uniqueness, given existence, needs only the elementary fact above. A map that is retroazimuthal and conformal only as a *multi-valued* correspondence (a genuine covering-space construction, wrapping around the punctured sphere more than once) is not excluded by this theorem's hypotheses, which require single-valuedness explicitly; section 7.2's zero-monodromy finding means no such covering is needed to achieve single-valuedness, but does not by itself rule out that a deliberately multi-valued construction could exist as a formally different object. This is recorded as a boundary of the theorem's scope, not a gap in its proof.

*What this does and does not settle.* Section 7.2's todobox posed two questions. This section closes the first in full: within the class of single-valued, orientation-preserving conformal retroazimuthal maps, the vector-alignment family is not a restrictive special case at all, and the construction is unique up to a uniform scale. The second question — Maurer's (1919) own specific uniqueness claim, and what exactly his unread 1919 paper proves or assumes — is untouched by this section, which settles the mathematical question on its own terms rather than the historical-attribution one; section 7.2's original todobox on that point stands exactly as it was.

## 7.6 Revising the volume's four-item programme

chapter 1's list of four impossibility claims should accordingly be read with item (iv) corrected: not an impossibility result, but a compatibility result whose only obstruction is the topological one of chapter 6, already established independently. The volume's genuine four results are therefore: (i) no simultaneous conformal/equal-area map (chapter 4); (ii) Milnor's sharp global distortion bound (chapter 5); (iii) no global retroazimuthal direction field, by the hairy ball theorem (chapter 6); and, in place of the original (iv), the finding of this chapter that conformality adds *no further* obstruction beyond (iii) once a centre and a punctured domain are fixed. chapter 1 is updated to reflect this in the corresponding revision pass.

# Chapter 8

# The Equidistance Theorem: How Much Distance Can Ever Be Preserved

## 8.1 Why this theorem has no chapter yet

Every classical family of projections examined so far in this volume carries a name that promises one property and silently concedes the others: *conformal* projections promise angles, *equal-area* projections promise area, and chapter 4 proves the two promises cannot be kept together. What has not yet been asked directly is the question a beginner asks first: why is there no projection simply called "equidistant," full stop, promising to preserve *distance*? The cartographic literature has an answer already embedded in its naming conventions — "azimuthal equidistant," "conic equidistant," "plate carrée (equidistant cylindrical)" — every use of the word is qualified, either to a single center point or to a single family of lines (meridians, or one standard parallel). This chapter states and proves, as the fifth member of the family standing beside chapters 4 to 7, why the qualification is never removable, and then asks the sharper question the qualification invites: exactly how much distance-exactness *is* achievable, and from how many points at once.

**Epistemic Notice**

The impossibility half of this chapter (theorem 8.2) is a direct corollary of theorem 2.1, proved from scratch in chapter 2; no new citation is needed for it. The achievability half draws on the historical two-point equidistant projection of Maurer (1919) and Close (1921) and the three-point Chamberlin trimetric projection (1946); both are described here from the secondary account in Snyder, *Flattening the Earth* (1993) and the current ArcGIS Pro projection reference, not from a primary-verified reading of Maurer's or Close's own papers, which remain unobtained — flagged accordingly. The genericity argument of proposition 8.6 is original to this volume.

## 8.2 The impossibility of full equidistance

**Definition 8.1.** A map $f\colon U \to \mathbb{R}^2$, for $U$ an open subset of the sphere $S^2_R$ of radius $R$, is *distance-preserving* (or *isometric*) if for every pair of points $x, y \in U$, the Euclidean distance $|f(x) - f(y)|$ equals the great-circle distance $d_S(x, y)$.

**Theorem 8.2** (No open region admits a fully distance-preserving map)**.** *For every open $U \subseteq S^2_R$ with $U \neq \emptyset$, no distance-preserving map $f\colon U \to \mathbb{R}^2$ exists.*

*Proof.* A distance-preserving map preserves the length of every curve, hence is an isometry of the induced first fundamental forms: it identifies the metric of $U$ with the flat Euclidean metric $du^2 + dv^2$. By theorem 2.1, Gaussian curvature is an isometry invariant, so $U$ would have to carry Gaussian curvature identically 0. But $S^2_R$ has constant Gaussian curvature $K = 1/R^2 \neq 0$ at every point (this is computed directly in chapter 2 from the standard spherical parametrization). Contradiction. □

This is the sense in which the theorem is a corollary rather than new mathematics: the sphere's curvature obstruction, once established in chapter 2, already rules out exact distance preservation *everywhere*, not merely between some particular pair of families of curves. The reason this deserves its own chapter, rather than a remark appended to chapter 2, is that — unlike conformality or equal-area, which are properties a projection either has or lacks outright over its whole domain — distance-exactness is naturally a *local and partial* property: a projection can be exact from a point, along a curve, or between a finite set of points, without being exact everywhere. The interesting mathematics is in measuring exactly how much of that partial exactness survives theorem 8.2's prohibition.

## 8.3 Distortion-free loci and Tissot's indicatrix

Recall from chapter 3 that at each point $p \in U$, a projection $f$ has a well-defined principal scale pair $(a(p), b(p))$, the semi-axes of the Tissot indicatrix at $p$, with $a(p) = b(p) = 1$ exactly when $f$ is locally an isometry at $p$ in every direction — what chapter 3 calls a *distortion-free point*.

**Proposition 8.3.** *The set of distortion-free points of any projection $f\colon U \to \mathbb{R}^2$ has empty interior in $U$.*

*Proof.* If the distortion-free set contained a nonempty open subset $V \subseteq U$, then $f$ restricted to $V$ would satisfy $a \equiv b \equiv 1$ on $V$, i.e. would be a local isometry on the open set $V$, contradicting theorem 8.2 applied to $V$. □

This is already enough to explain the qualifiers in every classical equidistant projection's name: none of them can be distortion-free on an open set, so each instead achieves distortion-freedom only along a lower-dimensional locus, and the projections differ precisely in *which* locus they choose.

**Proposition 8.4** (Azimuthal equidistant: exactness from a single point)**.** *Let $p_0 \in S^2_R$. The azimuthal equidistant projection centered at $p_0$, defined in polar coordinates $(\rho, \theta)$ on the plane by $\rho = Rc$ and $\theta = \alpha$, where $c$ is the angular great-circle distance from $p_0$ and $\alpha$ is the azimuth at $p_0$, satisfies $|f(x) - f(p_0)| = d_S(x, p_0)$ for every $x \in S^2_R \setminus \{-p_0\}$.*

*Proof.* By construction, $|f(x) - f(p_0)| = \rho(x) = Rc(x)$, and $Rc(x)$ is exactly the great-circle distance from $p_0$ to $x$ by definition of angular distance on a sphere of radius $R$. □

The azimuthal equidistant projection is thus distortion-free only *at* $p_0$ itself (a single point, consistent with proposition 8.3): the radial scale is exact everywhere by proposition 8.4, but the transverse scale away from $p_0$ is $Rc/\sin c \cdot 1/R = c/\sin c \neq 1$ for $c \neq 0$, so the indicatrix is a circle only in the limit $c \to 0$.

## 8.4 How many points can be simultaneously exact

The natural next question is whether a cleverer projection can push exactness from a single point to several points at once, and how far this can go. The answer is a genuine three-way split, and it is worth stating precisely because it is easy to guess wrong in both directions.

**Key Result**

*Exact simultaneous distance-preservation from two fixed points to every other point on the sphere is achievable (the two-point equidistant projection). Exact simultaneous distance-preservation from three or more fixed points, for a generic configuration of the points, is not achievable by any projection.*

### 8.4.1 Two points: achievable

**Proposition 8.5** (Two-point exactness)**.** *For any two distinct points $p_1, p_2 \in S^2_R$ with $d_S(p_1, p_2) < \pi R$, there is a map $f\colon S^2_R \setminus \Sigma \to \mathbb{R}^2$, for $\Sigma$ a suitable removed set of measure zero, such that $|f(x) - f(p_1)| = d_S(x, p_1)$ and $|f(x) - f(p_2)| = d_S(x, p_2)$ for every $x$ in the domain.*

We do not reprove the classical construction of Maurer (1919) and Close (1921) in full here — it proceeds by placing $f(p_1), f(p_2)$ at the two foci of a family of confocal ellipses and hyperbolas in the plane and mapping each point $x$ to the intersection of the ellipse and hyperbola carrying the correct pair of focal distances, a construction that is spelled out in Snyder (1993) and is a real, currently used projection (the National Geographic Society's atlas maps of Asia since 1959, among others) — but the geometric idea is worth stating because it is exactly what fails at three points.

**Todo — verify primary source**

Maurer's 1919 paper (*Annalen der Hydrographie und Maritimen Meteorologie* 47) and Close's independent 1921 note (*The Geographical Journal* 57(6), 446–448) have not been directly consulted for this volume; the construction above is reported from Snyder (1993) and current cartographic reference documentation. A later pass, recorded in appendix A.8.1, located a genuinely open (public-domain, no-paywall) digitization of the Close note and confirmed its precise title, byline, and grouping from the primary volume's own table of contents, but did not extract or read its three pages of mathematical content. A still later pass, recorded in appendix A.8.2, identified a plausible open-access digitization route for Maurer's paper as well (the BSH's own institutional digital library, which holds *Annalen* volumes from the nineteenth century onward with a page-addressable full-text API) but did not locate the specific 1919 volume 47 record within that collection, so Maurer's own text remains unread.

### 8.4.2 Three points: generically impossible

**Proposition 8.6** (Three-point exactness fails generically)**.** *Let $p_1, p_2, p_3 \in S^2_R$ be three points in general position (no two coincident, not all three on a common great circle). For a generic fourth point $x \in S^2_R$, there is no assignment of planar points $f(p_1), f(p_2), f(p_3), f(x)$ — with $f(p_1), f(p_2), f(p_3)$ fixed once and for all, independent of $x$ — satisfying all three distance conditions $|f(x) - f(p_i)| = d_S(x, p_i)$, $i = 1, 2, 3$, simultaneously.*

*Proof.* Fix planar images $f(p_1), f(p_2), f(p_3)$ forming a triangle with side lengths $d_S(p_1, p_2)$, $d_S(p_2, p_3)$, $d_S(p_1, p_3)$ — this is possible, and unique up to a rigid motion and a reflection, exactly because these three numbers individually satisfy the planar triangle inequality (a consequence of the corresponding spherical triangle inequality, which the three fixed points satisfy trivially). Now consider the requirement on $f(x)$ from just the first two conditions:

$|f(x) - f(p_1)| = d_S(x, p_1)$ and $|f(x) - f(p_2)| = d_S(x, p_2)$ place $f(x)$ at one of at most two points, the intersection of two circles of the given radii centered at $f(p_1)$ and $f(p_2)$ (generically exactly two, reflections of each other across the line $f(p_1)f(p_2)$, by the same construction underlying proposition 8.5). The third condition, $|f(x) - f(p_3)| = d_S(x, p_3)$, is then a single additional scalar equation imposed on a configuration — the two candidate points and the fixed third radius $d_S(x, p_3)$ — that already has no free parameter left to satisfy it with. Whether one, both, or neither of the two candidate points satisfies the third distance is now a closed, generically transverse condition on $x$'s position (equivalently, on the triple of spherical distances $(d_S(x, p_1), d_S(x, p_2), d_S(x, p_3))$, which must satisfy the planar Cayley–Menger consistency relation for the quadruple $f(p_1), f(p_2), f(p_3), f(x)$ to close up in the plane at all): it holds on a codimension-one subset of $x \in S^2_R$ at best, never on an open set. Since $S^2_R$ is two-dimensional, a generic $x$ fails it. □

*Remark* 8.7. This is exactly the mathematical content behind the existence of the *Chamberlin trimetric* projection (1946), the natural three-point extension of proposition 8.5: it does not achieve exact three-point equidistance, because proposition 8.6 shows no projection can, but instead places $f(x)$ at (approximately) the centroid of the triangle formed by the three candidate points obtained from each pair of the three distance constraints separately, minimizing rather than eliminating the residual error. Waldo Tobler's later generalization to an arbitrary number of control points, cited in the same secondary literature, replaces exact construction with root-mean-square minimization for exactly this reason.

## 8.5 Where this leaves the taxonomy

Chapter 26 classifies named projections along the axes of conformality, equal-area, and interruption. Propositions 8.3 to 8.6 together describe a genuinely separate axis — the size and shape of the exact-distance locus, which for real projections in use is empty (Mercator), a single point (azimuthal equidistant), a pair of points (two-point equidistant), or, in the trimetric and higher-order cases, provably never exact but only minimized. Folding this fifth axis formally into chapter 26's table is left as an explicit open item for that chapter rather than attempted here, to avoid extending an already carefully closed chapter on the strength of a single new addition.

# Chapter 9

# Beyond the Sphere: Why the Impossibility Is the Sphere's, Not Every Surface's

## 9.1 A question this volume has so far begged

Every impossibility theorem proved so far — theorems 4.1 and 8.2 (see chapter 4), the hairy ball obstruction of chapter 6, the retroazimuthal-conformal exclusion of chapter 7 — is stated for the sphere $S^2$, because the sphere is the shape of the Earth and this is a volume about cartography. But every one of the proofs, on inspection, uses a specific fact about the sphere that is not true of surfaces in general: that its Gaussian curvature is everywhere positive (chapter 2), that its Euler characteristic is $\chi(S^2) = 2$ (used via Poincaré–Hopf in chapter 6 and again, for the icosahedral discretization, in chapter 25). This chapter asks the question the rest of the volume has left implicit: is cartographic impossibility a fact about closed surfaces in general, or specifically about the sphere? The answer is genuinely surprising the first time one sees it, and it sharpens rather than dilutes the volume's central claim.

**Epistemic Notice**

This chapter is self-contained: every result either restates material already proved in chapters 2 and 6 in more general form, or is a standard fact of differential topology (the Gauss–Bonnet theorem for closed surfaces, and the classification of closed orientable surfaces by genus) stated here without independent primary-source verification, flagged accordingly. No new cartographic case study is claimed; the point of the chapter is strictly to locate the sphere correctly within the wider family.

## 9.2 The Gauss–Bonnet theorem for closed surfaces

**Theorem 9.1** (Gauss–Bonnet, global form)**.** *Let $\Sigma$ be a smooth, closed (compact, boundaryless), orientable surface with Gaussian curvature $K$ and Euler characteristic $\chi(\Sigma)$. Then*

$$\iint_{\Sigma} K \, dA = 2\pi\chi(\Sigma).$$

*Proof.* The proof rests on one classical building block, stated here without independent rederivation and flagged for that reason alone: the *local* Gauss–Bonnet formula for a geodesic triangle $T$ with interior angles $\alpha, \beta, \gamma$,

$$\iint_{T} K \, dA \;=\; (\alpha + \beta + \gamma) - \pi,$$

which holds because $T$'s three sides are geodesics (so their geodesic curvature is identically zero) and the general local Gauss–Bonnet identity for a simply connected region bounded by a piecewise-geodesic curve, $\iint_T K\,dA + \oint_{\partial T} k_g\,ds + \sum(\text{exterior angles}) = 2\pi$, then reduces to the displayed formula once $k_g \equiv 0$ and each exterior angle $\pi - \alpha$ etc. is substituted. Given this, the global theorem is a fully self-contained combinatorial argument:

*Step 1 (triangulation).* Every smooth closed surface admits a triangulation by finitely many geodesic triangles, fine enough that each triangle lies in a normal-coordinate neighbourhood of one of its vertices (existence of such a triangulation is itself a standard fact of differential topology, not re-derived here). Write $V$, $E$, $F$ for the resulting vertex, edge, and face counts; since the surface is closed (no boundary) and every face is a triangle, each edge borders exactly two faces and each face has exactly three edges, so counting edge-face incidences two ways gives

$$3F = 2E.$$

*Step 2 (sum the local formula over faces).* Summing the local Gauss–Bonnet formula over all $F$ triangles,

$$\iint_\Sigma K\,dA \;=\; \sum_{\text{faces}} \iint_T K\,dA \;=\; S - F\pi,$$

where $S$ denotes the sum, over every triangle, of that triangle's three interior angles (so $S$ is the grand total of all angles of all triangles, each counted once).

*Step 3 (regroup the same sum by vertex).* $S$ can equally be computed by grouping angles by which vertex they sit at rather than which triangle they belong to: $S = \sum_v(\text{sum of the angles of the triangles meeting at } v)$. At each vertex $v$ of a triangulation of a closed surface, the triangles meeting at $v$ form a closed fan that tiles a full neighbourhood of $v$ without gap or overlap — this is a statement about how triangles fit together combinatorially at an interior point of a triangulated 2-manifold, not a curvature statement, and holds regardless of what $K$ does nearby (all of the curvature's effect on the total is already accounted for face-by-face in Step 2) — so the angles at each vertex sum to exactly $2\pi$. Hence $S = 2\pi V$.

*Step 4 (combine).* Substituting $S = 2\pi V$ into Step 2's identity,

$$\iint_\Sigma K\,dA \;=\; 2\pi V - \pi F.$$

Using Step 1's edge count $E = \frac{3}{2}F$, i.e. $2\pi E = 3\pi F$, so $\pi F = 2\pi E - 2\pi F$, and substituting:

$$2\pi V - \pi F \;=\; 2\pi V - (2\pi E - 2\pi F) \;=\; 2\pi(V - E + F) \;=\; 2\pi\chi(\Sigma),$$

which is exactly theorem 9.1's claim. □

The full triangulation argument above (Steps 1–4) is now derived explicitly in this volume rather than merely cited, reducing the theorem to a single remaining classical input: the local Gauss–Bonnet formula for a single geodesic triangle. That local formula — and, one level further back, the existence of a fine enough geodesic triangulation of an arbitrary smooth closed surface — are genuinely standard differential-topological facts (available in do Carmo or Spivak) that this volume has not re-derived from a primary 19th-century source, and are flagged on that basis, consistent with chapter 2's own practice of distinguishing what is derived here from what is assumed. This is a narrower and more honest flag than the previous todobox's blanket "standard material, not re-derived," which covered the triangulation-to-Euler-characteristic argument as well; that part is no longer merely cited.

Closed orientable surfaces are classified, up to homeomorphism, by a single nonnegative integer $g$, the *genus* (informally, the number of handles): $\chi(\Sigma_g) = 2 - 2g$. The sphere is $g = 0$, with $\chi = 2$; the torus is $g = 1$, with $\chi = 0$; the double torus is $g = 2$, with $\chi = -2$; and so on.

## 9.3 What Gauss–Bonnet forces at each genus

**Corollary 9.2.** *On $\Sigma_0 = S^2$, $\iint K\,dA = 4\pi > 0$: the total curvature is forced strictly positive, so $K$ cannot be identically zero, and (by continuity, since $\Sigma_0$ is connected) $K$ must be strictly positive on some open set. This is consistent with, but weaker than, the sphere's actual constant positive curvature used throughout chapters 2 and 3 and the rest of this volume — theorem 9.1 alone only forces* some *positive curvature somewhere, not the round metric's curvature everywhere.*

**Corollary 9.3.** *On $\Sigma_1$ (any surface of genus 1, topologically a torus), $\iint K\,dA = 0$.*

Corollary 9.3 is where the surprise sits. Total curvature zero does *not* by itself mean $K \equiv 0$ pointwise — an embedded doughnut-shaped torus in $\mathbb{R}^3$, for instance, has strictly positive curvature on its outer equator and strictly negative curvature on its inner equator, integrating to zero overall, and is very much not developable. But total curvature zero is *consistent with* $K \equiv 0$ everywhere, in a way that is never consistent for the sphere by corollary 9.2 — and genus-1 surfaces are exactly the case where this potential is realized.

**Theorem 9.4** (The flat torus exists)**.** *There exists a Riemannian metric on the torus $T^2 = \mathbb{R}^2/\Lambda$, for $\Lambda$ any rank-2 lattice in $\mathbb{R}^2$, with Gaussian curvature identically zero everywhere.*

*Proof.* Equip $\mathbb{R}^2$ with the standard flat metric $du^2 + dv^2$, which has $K \equiv 0$ by direct computation from eq. (2.1) (or trivially, since $E = G = 1, F = 0$ are constant, every term in the Brioschi/Codazzi–Mainardi expression from chapter 2 vanishes). This metric is invariant under the translation action of $\Lambda$, hence descends to a well-defined smooth Riemannian metric on the quotient $T^2 = \mathbb{R}^2/\Lambda$, still with $K \equiv 0$ at every point, since curvature is computed from local data and the quotient map is a local isometry. □

**Key Result**

*The impossibility results of this volume — no global isometric flattening (theorem 8.2), no simultaneous conformal-and-equal-area map (chapter 4) — are theorems* about the sphere specifically*, not about closed surfaces in general. They follow from $\chi(S^2) = 2 \neq 0$ forcing strictly positive curvature somewhere (ruling out flatness) and, for the sphere's particular constant-curvature metric, everywhere. The torus, with $\chi = 0$, has no such obstruction: it admits a perfectly flat metric, and a region of the flat torus* can *be isometrically unrolled onto a flat sheet of paper (this is literally how a video-game world with wraparound edges, or a physical torus built by gluing opposite edges of a flat rectangle, is constructed).*

## 9.4 Why this does not touch cartography, and why it should still be in this volume

One might ask why this matters for a volume about maps of the Earth, which is a sphere (to excellent approximation — see chapters 16 and 17 for the ellipsoidal correction) and will never be a torus. The answer is that it sharpens exactly what this volume has been claiming all along. Theorems 4.1 and 8.2 are not consequences of "surfaces resist flattening" as a vague slogan; a surface's resistance to flattening is entirely a statement about its Euler characteristic

forcing nonzero total curvature, via theorem 9.1, and the sphere is the topologically simplest case, $\chi = 2$, where that force is at its *strongest* — every other closed orientable surface has smaller $|\chi|$ (down to exactly 0 at the torus, where the force vanishes entirely) or, for $g \geq 2$, a total curvature constraint of the *opposite* sign, forcing negative curvature somewhere instead of positive. The Earth's cartographic impossibility is, in this precise sense, closer to a maximal case than a generic one: no closed surface is harder to flatten, topologically, than the sphere is.

*Remark* 9.5. This also gives a one-line reason chapter 25's discrete computation was bound to find total angular defect $4\pi$ (equivalently 720°) for every triangulated approximation to the sphere, regardless of how the icosahedral refinement's individual vertex defects were distributed among pentagons and hexagons or which of the three triangulation conventions of section 25.7.48 was used: the discrete Gauss–Bonnet identity $\sum_v \delta_v = 2\pi\chi$ is the same $\chi = 2$ forcing $\iint K\, dA = 4\pi$ in the continuous case, and no amount of refinement or choice of convention can move a topological invariant.

# Chapter 10

# The Negative Case, and Why Conformality Never Needed the Sphere at All

## 10.1 Completing a trichotomy, cautiously

Chapter 9 showed that closed orientable surfaces split by genus into a case Gauss–Bonnet forces toward positive total curvature ($g = 0$, the sphere), a case it leaves exactly balanced at zero ($g = 1$, the torus, which achieves a flat metric), and — mentioned there only in a closing remark — a case ($g \geq 2$) it forces toward negative total curvature. This chapter is explicitly a completion of that remark, not an independent structural gap: it exists because leaving the trichotomy positive/zero/*unstated* after devoting a full chapter to the first two cases would be a strange asymmetry, not because cartography needs negatively curved surfaces. It is filed immediately after chapter 9 on that basis.

**Epistemic Notice**

The existence of constant-curvature hyperbolic metrics on every surface of genus $g \geq 2$ (theorem 10.3 below) is the Poincaré–Koebe uniformization theorem, one of the deepest results in the subject; it is stated here, not proved, and flagged as such — unlike chapter 9's theorem 9.4, which this volume does prove from scratch, the hyperbolic existence statement is well beyond what this volume can re-derive and is reported from the standard literature (e.g. Farkas–Kra, *Riemann Surfaces*) without primary verification.

## 10.2 The hyperbolic plane

**Definition 10.1.** The *hyperbolic plane* $\mathbb{H}^2$ is the open unit disk $\{(u, v) : u^2 + v^2 < 1\}$ equipped with the Poincaré metric

$$ds^2 = \frac{4(du^2 + dv^2)}{(1 - u^2 - v^2)^2}.$$

**Proposition 10.2.** *The Gaussian curvature of* $\mathbb{H}^2$ *is identically* $K \equiv -1$.

*Proof sketch.* The metric is conformal, $ds^2 = e^{2\sigma}(du^2 + dv^2)$ with $\sigma(u, v) = \ln 2 - \ln(1 - u^2 - v^2)$. For a conformal metric the standard formula (a direct consequence of the same Codazzi–Mainardi machinery used in chapter 2, specialized to $F = 0, E = G = e^{2\sigma}$, and not re-derived here in full to avoid duplicating chapter 2's general computation) gives $K = -e^{-2\sigma}\Delta\sigma$, where $\Delta$ is the flat Laplacian. Computing $\Delta\sigma$ for $\sigma$ as above and substituting gives $K \equiv -1$ after

simplification; this is a standard computation reported, not independently re-verified line by line, from Farkas–Kra. □

By the same logic as corollaries 9.2 and 17.2: since $K \equiv -1 \neq 0$ everywhere on $\mathbb{H}^2$, theorem 2.1 again forbids any open region of $\mathbb{H}^2$ from being isometrically flattened onto the ordinary Euclidean plane. Negative curvature obstructs isometry exactly as positive curvature does — theorem 2.1 cares only that $K \neq 0$, not its sign, a fact already implicit in chapter 2's statement but worth making explicit here since it is easy to guess, wrongly, that negative curvature is somehow "the opposite" problem rather than the same one.

## 10.3 What genus $g \geq 2$ gets, that the sphere and torus do not

**Theorem 10.3** (Uniformization, hyperbolic case)**.** *Every closed orientable surface $\Sigma_g$ of genus $g \geq 2$ admits a Riemannian metric of constant Gaussian curvature $K \equiv -1$.*

**Corollary 10.4.** *Every closed orientable surface admits a constant-curvature metric, and the sign of that curvature is determined by genus alone: positive for $g = 0$ (the round sphere, curvature $+1/R^2$, this volume's subject throughout), zero for $g = 1$ (the flat torus of theorem 9.4), negative for $g \geq 2$ (theorem 10.3). Genus and curvature sign are in exact correspondence via $\chi = 2 - 2g$ and theorem 9.1: $\chi > 0 \Leftrightarrow g = 0$, $\chi = 0 \Leftrightarrow g = 1$, $\chi < 0 \Leftrightarrow g \geq 2$.*

## 10.4 Why conformal maps exist regardless, and what this says about the whole volume

Everything so far in this chapter has been about *isometry* (distance preservation), which theorem 2.1 obstructs whenever $K \neq 0$ — true for the sphere, true for $\mathbb{H}^2$, true for every hyperbolic $\Sigma_g$. But this volume's other central theorem, chapter 4's result that conformality and equal-area cannot coexist, is a much narrower claim than "conformal maps to the plane don't exist": chapters 14 and 16 construct working conformal maps of the sphere at every step. The natural question raised by corollary 10.4, and the reason this chapter earns its place rather than merely restating chapter 9, is whether conformality is available this freely regardless of which of the three curvature signs a surface has.

**Key Result**

*[Uniformization, full statement] Every simply connected Riemann surface is conformally equivalent to exactly one of three model spaces: the round sphere $S^2$ (curvature $+1$), the Euclidean plane $\mathbb{C}$ (curvature $0$), or the hyperbolic disk $\mathbb{H}^2$ (curvature $-1$). Consequently, the universal cover of* any *smooth surface admits a conformal identification with one of these three models, regardless of the surface's own actual curvature: a small enough conformal chart exists locally near every point of every surface, positively curved, flat, or negatively curved alike.*

**Todo — verify primary source**

Section 10.4 is the full Poincaré–Koebe uniformization theorem in its classical statement; it is quoted, not proved, and no primary 1907 source (Poincaré; Koebe, independently) has been consulted for this volume. A later pass identified precise citations for both — Poincaré, "Sur l'uniformisation des fonctions analytiques," *Acta Mathematica* 31, pp. 1–63 (received 1906, published 1907/1908), and Koebe's first of three 1907–1908 communications, "Über die Uniformisierung beliebiger analytischer Kurven," *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1907, pp. 191–210 (with a second communication at pp. 633–669 of the same 1907 volume and a third at pp. 337–358 of the 1908 volume) — and confirmed, via International Press's public statement on Acta Mathematica's own journal page, that the complete historical Acta Mathematica archive back to 1882 is provided in fully open access by special arrangement with the Institut Mittag-Leffler. Fetching the article itself via Project Euclid's hosting was blocked in this pass by an automated bot-protection challenge (Incapsula) rather than by any subscription or paywall notice, so this remains an access-confirmed-but-text-unread item in the same sense as chapter 22's Christoffel/Schwarz pair and chapter 16's Krüger monograph, not a genuinely closed one. A later pass resolved the small page-range discrepancy left open above (some secondary sources gave Poincaré's closing page as 64 rather than 63): the complete, independently-compiled table of contents of *Acta Mathematica* maintained at German Wikisource records Band 31 (1908) as running *H. Poincaré: Sur l'uniformisation des fonctions analytiques, S. 1–63*, a third source in agreement with the citation already adopted here and with Snyder/Farkas–Kra, against the lone outlier giving 64. This closes the page-range question without requiring the primary text itself, which remains unread. Neither Poincaré's nor Koebe's own derivation has been checked against theorem 10.3 and section 10.4's statements as given here.

This is the sentence that ties chapters 9 and 10 back to the volume's actual subject. Theorem 2.1's obstruction — the one this entire volume is built on — is specifically an obstruction to *isometry*, and it is exactly as strict for $K < 0$ as for $K > 0$: no curvature sign gets a pass. Section 10.4's obstruction-free existence of *conformal* charts, by contrast, holds for every curvature sign equally, with no exceptions at all — which is precisely why chapters 14, 16 and 17 were always able to construct a working conformal projection, on the sphere or the ellipsoid, without needing to fight any topological or curvature obstruction for that half of the problem. The obstruction chapter 4 proves is not "a conformal map cannot exist": it is the strictly narrower, and in this light much more interesting, claim that a conformal map cannot *also* be equal-area. Conformality alone was never the hard part; chapter 4 is a theorem about the interaction of two properties, not about either one in isolation, and section 10.4 is the reason the volume never had to prove a separate "conformal maps exist" theorem before combining it with equal-area to get an impossibility — existence was free the entire time.

# Chapter 11

# Hilbert's Theorem: The Hyperbolic Plane Cannot Even Be Built

## 11.1 A question chapter 10 left unasked

Chapter 10 established that $\mathbb{H}^2$'s curvature $K \equiv -1$ forbids, by theorem 2.1, any isometry between an open piece of $\mathbb{H}^2$ and the flat Euclidean plane — exactly the same obstruction, and for exactly the same reason, that forbids a flat map of any piece of the sphere. That symmetry was the chapter's own closing point: theorem 2.1 does not care about the sign of $K$, only that $K \neq 0$.

This chapter asks a different question, one chapter 10 did not raise because it falls outside this volume's usual subject of maps *to the printed page*. The sphere's curvature obstruction to a flat map does not stop the sphere from existing as an actual surface in ordinary three-dimensional space: a terrestrial globe is a literal, complete, unbounded-in-no-direction physical realization of a surface of constant positive curvature, sitting in $\mathbb{R}^3$ with its own intrinsic metric matching the abstract round sphere's exactly. The natural question this raises is whether $\mathbb{H}^2$ admits an analogous physical realization — not a flat map, but a literal curved surface in $\mathbb{R}^3$, complete in every direction, whose intrinsic geometry is the hyperbolic plane's. The answer, due to Hilbert in 1901, is no, and the reason is a different theorem from theorem 2.1, resting on a global argument theorem 2.1's own local, pointwise mechanism cannot supply.

**Epistemic Notice**

This chapter's central theorem (theorem 11.8 below) was originally stated and sketched from the standard textbook treatment of the subject (do Carmo, *Differential Geometry of Curves and Surfaces*), not from Hilbert's own 1901 paper. A later pass has since located and read that paper in full (appendix A's appendix A.9); section 11.5's proof sketch below now follows Hilbert's own argument directly, and section 11.6 below adds a second theorem from the same paper not previously reported in this volume. What remains a *sketch*, even against the primary source, is stated precisely at each step rather than left general: this chapter reports Hilbert's logical skeleton and the two area computations that drive his contradiction, without independently re-proving every one of his four asymptotic-curve lemmas line by line. The two supporting computations — the pseudosphere's curvature, and the sine-Gordon identity underlying its asymptotic-coordinate description — remain verified independently in this pass by direct symbolic computation (sections 11.2 and 11.3 below), not merely asserted from any source.

## 11.2 Beltrami's pseudosphere: a local realization, and where it fails

**Definition 11.1** (Tractrix and pseudosphere)**.** The *tractrix* is the plane curve traced by a point dragged along a straight line by an inextensible string of fixed length $a$ attached to a puller moving along a perpendicular line, parametrized for $u \in (0, \pi/2)$ by

$$x(u) = a\left(\sin u,\ \cos u + \ln\tan\tfrac{u}{2}\right).$$

The *pseudosphere* (or *tractricoid*) of pseudo-radius $a$ is the surface of revolution obtained by rotating this curve about the line $x = 0$ (its asymptote), parametrized for $u \in (0, \pi/2)$, $v \in (0, 2\pi)$ by

$$X(u, v) = a\big(\sin u\cos v,\ \sin u\sin v,\ \cos u + \ln\tan\tfrac{u}{2}\big).$$

**Proposition 11.2.** *The pseudosphere has constant Gaussian curvature $K \equiv -1/a^2$ at every regular point.*

**Key Result**

*[Verified by direct symbolic computation, not merely cited] Proposition 11.2 is confirmed by* `scripts/hilbert_pseudosphere_curvature.py`*, which computes the first fundamental form $(E, F, G)$ and second fundamental form $(e, f, g)$ of $X(u, v)$ directly from the parametrization above, forms $K = (eg - f^2)/(EG - F^2)$ symbolically, and evaluates the resulting (unsimplified — sympy's automatic simplification stalls on the $\tan(u/2)$ term's branch ambiguities) expression numerically to 30 digits of precision at five test points spanning $a \in \{0.5, 1, 2, 3\}$ and $u, v$ ranging across the parameter domain. All five evaluations match $-1/a^2$ to within $9 \times 10^{-16}$, floating-point roundoff at that precision. The script additionally confirms $F \equiv 0$ (an orthogonal parametrization) and that $E, G$ do not depend on $v$ (the surface-of-revolution symmetry), both sanity checks on the parametrization itself rather than on the curvature claim.*

The pseudosphere is therefore a genuine, if partial, realization of $\mathbb{H}^2$'s geometry: a literal surface in $\mathbb{R}^3$ whose intrinsic metric has constant curvature $-1/a^2$, matching $\mathbb{H}^2$'s own curvature $-1$ exactly when $a = 1$. It is the negative-curvature analogue of the sphere in exactly the sense section 11.1 above was asking for — with one immediate and visible defect.

*Remark* 11.3 (The pseudosphere is not complete)*.* The parametrization above is only defined for $u \in (0, \pi/2)$: as $u \to \pi/2^-$, $\ln\tan(u/2) \to \ln 1 = 0$ and the surface closes up smoothly along its equator of radius $a$, but as $u \to 0^+$, $\ln\tan(u/2) \to -\infty$ while $\sin u \to 0$, so the surface pinches to a sharp, non-smooth cuspidal edge at a finite distance along every generating tractrix. A geodesic on the pseudosphere running toward this edge reaches it in finite arc length and cannot be continued past it as a $C^2$ curve on the surface: the pseudosphere is a surface with boundary, in the sense that matters for completeness, even though the parametrization has no boundary curve written into its formula explicitly. It is, in the language chapter 10 borrowed from Riemannian geometry, an *incomplete* realization of constant curvature $-1$: a bounded patch, not the whole of $\mathbb{H}^2$.

This is not a defect of this particular parametrization that a cleverer one might fix. It is, as theorem 11.8 below shows, unavoidable: no surface of constant negative curvature in $\mathbb{R}^3$ can be completed past its cuspidal edges, no matter how it is built.

## 11.3 The asymptotic net, and the sine-Gordon equation

The proof of theorem 11.8 works not with the pseudosphere's surface-of-revolution parametrization but with a different coordinate system adapted specifically to negative curvature, which this section introduces because section 11.5 below needs it and because it

produces, as a byproduct, a second and independent confirmation that the pseudosphere realizes $K = -1$.

**Definition 11.4** (Asymptotic curves)**.** On a surface with $K < 0$ everywhere, at each point the second fundamental form has two real, distinct null directions (directions along which normal curvature vanishes); the two families of integral curves of these directions are the surface's *asymptotic curves*. Coordinates $(u, v)$ aligned with the two asymptotic families at every point are *asymptotic coordinates*, and in such coordinates the first fundamental form takes the special orthogonal-in-neither-direction form

$$ds^2 = du^2 + 2\cos\omega\, du\, dv + dv^2,$$

where $\omega(u, v) \in (0, \pi)$ is the angle between the two asymptotic directions at $(u, v)$.

**Proposition 11.5** (Gauss equation in asymptotic coordinates)**.** *For a surface parametrized by asymptotic coordinates as above, the Gauss equation relating the metric to the curvature reduces to the* sine-Gordon equation

$$\omega_{uv} = K \sin\omega.$$

*In particular, for a surface of constant curvature $K \equiv -1$, $\omega$ satisfies $\omega_{uv} = -\sin\omega$ if $\omega$ is measured as the angle to the* negative *asymptotic direction, or, in the sign convention this chapter adopts (matching the standard treatment), $\omega_{uv} = \sin\omega$.*

**Key Result**

*[The pseudosphere's asymptotic angle function, exhibited and independently verified] The function*

$$\omega(u, v) = 4\arctan\left(e^{u+v}\right)$$

*solves $\omega_{uv} = \sin\omega$ identically, confirmed by direct symbolic differentiation in* `scripts/hilbert_sine_gordon_check.py`*, which computes $\omega_{uv} - \sin\omega$ symbolically and simplifies it to exactly $0$. Three superficially similar candidates — with the sign of $v$ flipped, with a different scaling of $u + v$, and a reflected form — are checked alongside it in the same script and confirmed, equally directly,* not *to solve the equation, establishing that the match is not an artifact of an overly permissive simplification but a genuine, isolated solution among nearby alternatives. This is the standard "static kink" (one-soliton) solution of the sine-Gordon equation, and its geometric meaning is exactly the pseudosphere: restricted to the domain $u + v < 0$, $\omega$ increases monotonically from $0$ (as $u + v \to -\infty$) toward $\pi$ (as $u + v \to 0^-$), and the metric's degeneration as $\cos\omega \to -1$ at this upper limit — the first fundamental form $du^2 + 2\cos\omega\, du\, dv + dv^2$ collapsing toward $(du - dv)^2$, a rank-deficient quadratic form — is the asymptotic-coordinate description of exactly the same cuspidal-edge incompleteness remark 11.3 identified directly from the surface-of-revolution parametrization. The two coordinate systems, arrived at independently in this chapter, agree on where the pseudosphere breaks down.*

**Todo — verify primary source**

Section 11.3's geometric reading of the domain $u + v < 0$ (and its identification of $\omega \to \pi$ with the pseudosphere's cuspidal edge specifically, as opposed to some other degeneracy) follows the standard account of the pseudosphere/sine-Gordon correspondence found in the Bäcklund-transformation literature (e.g. Rogers and Schief, *Bäcklund and Darboux Transformations*), not independently re-derived here from the surface-of-revolution parametrization by an explicit change of coordinates; only the algebraic fact that $\omega$ solves the sine-Gordon equation is independently verified by this pass's own symbolic computation. A future pass constructing the explicit $(u, v) \to$ tractricoid-parametrization coordinate change would close this gap outright.

## 11.4 Hazzidakis's area formula

The proof of theorem 11.8 needs one further classical fact, which this chapter states without independent re-derivation.

**Proposition 11.6** (Hazzidakis, 1880)**.** *Let $R$ be a simply connected region of a surface of constant curvature $K \equiv -1$, bounded by a quadrilateral whose four sides are arcs of asymptotic curves, with interior angles $\alpha, \beta, \gamma, \delta$ at its four vertices. Then the area of $R$ is*

$$\operatorname{Area}(R) = (\alpha + \beta + \gamma + \delta) - 2\pi.$$

*Remark* 11.7. Proposition 11.6 is a direct analogue, for asymptotic quadrilaterals on a $K \equiv -1$ surface, of the Gauss–Bonnet formula chapter 9 used for geodesic triangles: it converts an area integral of curvature into a boundary term built purely from angles. Its immediate and load-bearing consequence for section 11.5 below is that **the area of any such quadrilateral is strictly less than** $2\pi$, since each of the four angles is strictly less than $\pi$, so their sum is strictly less than $4\pi$.

**Key Result**

*[Derived from scratch and independently verified, not merely cited] Proposition 11.6 is derived here directly from proposition 11.5's sine-Gordon equation together with elementary corner-angle bookkeeping, rather than quoted from Hazzidakis's, Hilbert's, or do Carmo's own texts. Take $R$ to be the coordinate rectangle $[u_1, u_2] \times [v_1, v_2]$ of an asymptotic net. Since $ds^2 = du^2 + 2\cos\omega\, du\, dv + dv^2$ gives $\operatorname{Area}(R) = \iint_R \sin\omega\, du\, dv$ and proposition 11.5 gives $\sin\omega = \omega_{uv}$ (in this chapter's own verified sign convention, section 11.3), the mixed-partial fundamental theorem of calculus converts the area integral directly into corner values of $\omega$:*

$$\operatorname{Area}(R) = \iint_R \omega_{uv}\, du\, dv = \omega(u_2, v_2) - \omega(u_2, v_1) - \omega(u_1, v_2) + \omega(u_1, v_1).$$

*Separately, inspecting the outgoing tangent direction of each of $R$'s two edges at each of its four corners (a pure bookkeeping exercise, independent of the differential equation) shows the interior angle equals $\omega$ at the corners $(u_1, v_1)$ and $(u_2, v_2)$, where both edges leave along like-signed axes ($+X_u, +X_v$ or $-X_u, -X_v$), and equals $\pi - \omega$ at the other two corners $(u_2, v_1)$ and $(u_1, v_2)$, where the edges leave along opposite-signed axes. Summing all four gives $\alpha + \beta + \gamma + \delta = \big[\omega(u_1, v_1) - \omega(u_2, v_1) + \omega(u_2, v_2) - \omega(u_1, v_2)\big] + 2\pi$, which is exactly the area expression above, plus $2\pi$: $\operatorname{Area}(R) = (\alpha + \beta + \gamma + \delta) - 2\pi$, as stated.*
*This is confirmed two further ways by `scripts/hazzidakis_from_scratch.py`: symbolically, as an identity in the four corner values of $\omega$ treated as free symbols (no specific solution needed); and numerically, on three concrete rectangles strictly inside the pseudosphere's regular domain $u + v < 0$, using section 11.3's own kink solution $\omega(u, v) = 4\arctan(e^{u+v})$, by computing $\operatorname{Area}(R)$ two independent ways — direct 40-digit 2D quadrature of $\sin\omega$ over the rectangle, and the corner-angle formula above — and confirming agreement to within $10^{-40}$. The same script confirms, on the same three rectangles, that the mirror-image form $2\pi - (\text{angle sum})$ disagrees with the directly-quadrature-computed area by an amount equal to twice the true area itself (i.e. it has exactly the wrong sign, not a small error), settling numerically that the two candidate formulas are not both approximately right.*

**Todo — verify primary source**

The derivation above resolves this chapter's own gap in deriving proposition 11.6, and section 11.5 below records how it also resolves the Hilbert/Darboux sign discrepancy. What it does *not* touch is the historical attribution: proposition 11.6 is still reported here under Hazzidakis's name only from its standard citation in the differential-geometry textbook literature (do Carmo's treatment), attributing it to J. N. Hazzidakis's 1880 paper in *Crelle's Journal*; this pass has independently re-derived the formula itself but has not independently confirmed the precise volume and page range of Hazzidakis's paper, nor consulted the original, nor established that Hazzidakis's own proof matches the corner-angle argument given here rather than some other route. Appendix A's own record of this point (appendix A.9) notes the same open citation gap; a future pass locating Hazzidakis's 1880 paper directly, in the same spirit as this volume's other primary-source recovery passes, would close it.

## 11.5 Hilbert's theorem

**Theorem 11.8** (Hilbert, 1901)**.** *There is no complete, regular ($C^2$) surface of constant negative Gaussian curvature isometrically immersed in Euclidean space $\mathbb{R}^3$.*

*Proof sketch, following Hilbert's own 1901 argument.* Suppose, for contradiction, that $S \subset \mathbb{R}^3$ is such a surface, with $K \equiv -1$ (the general constant-negative-$K$ case rescales to this one) and complete in the sense that every geodesic of $S$ extends indefinitely. Fix a point $O \in S$ and set up asymptotic coordinates $(u, v)$ centred at $O$ as in section 11.3, so $u = 0$ and $v = 0$ are the two asymptotic curves through $O$; completeness is used, as before, to argue that neither can terminate at a finite parameter value.

Hilbert shows this local chart is in fact global — a bijection between all of $\mathbb{R}^2$ and all of $S$ — through four short lemmas about the two families of asymptotic curves, none of which needs the sine-Gordon equation directly:

(a) *No asymptotic curve closes up on itself.* If some asymptotic curve $a$ returned to its own starting point, marking off an equal fixed arc length $s$ along the transverse asymptotic curve at every point of $a$, always to the same side, produces a second closed asymptotic curve $b$; the strip between $a$ and $b$ then forms a closed asymptotic quadrilateral whose four angles sum to exactly $2\pi$ at the two junctions, forcing its area, by proposition 11.6, to vanish — impossible for a strip of positive width.

(b) *Two asymptotic curves through a point meet nowhere else.* If a curve $b$ of the transverse family met $a$ a second time, an infinite sequence of asymptotic quadrilaterals built from the successive crossing points would need angles simultaneously increasing without bound and confined to $(0, \pi)$ — impossible.

(c) *No asymptotic curve self-intersects.* A double point would let the two loop-branches through it serve as new coordinate axes; repeating (b)'s construction locally around the loop then produces two curves of the same family meeting twice, contradicting (b).

(d) *A fixed-arc-length offset of an asymptotic curve, along its transverse family, never meets the original.* If it did, the further branch through the meeting point would have to belong to the offset curve too, making that point a double point of the offset curve — forbidden by (c).

Together, (a)–(d) show the two families of asymptotic curves foliate $S$ without self-crossings within a family and with exactly one crossing between families through every point of the plane, so $(u, v) \mapsto S$, extended over all of $\mathbb{R}^2$ by completeness, is a bijection; $S$ is in particular simply connected.

With this global chart established, Hilbert computes $S$'s total area two independent ways and contradicts one against the other.

*Bound one, from the asymptotic net.* Let $R_n$ be the asymptotic quadrilateral with opposite corners $(n, n)$ and $(-n, -n)$. By proposition 11.6, $\text{Area}(R_n) = (\text{angle sum of } R_n) - 2\pi < 2\pi$ for every $n$, since all four angles are strictly less than $\pi$ and so their sum is strictly less than $4\pi$. As $n \to \infty$, $R_n$ exhausts $S$ by the bijectivity just established, so $\text{Area}(S) = \lim_{n\to\infty} \text{Area}(R_n) \leq 2\pi$ — a single fixed bound, with no need to track how $\omega$ behaves as the quadrilateral grows.

*Bound two, from geodesic circles.* Independently, mark off the same geodesic distance $r$ along every geodesic emanating from $O$; because $K \equiv -1$ makes each such geodesic segment shortest between its endpoints, two of them cannot cross again without one ceasing to be shortest beyond the crossing, so the endpoints trace a simple closed curve. The region it bounds is, intrinsically, an ordinary hyperbolic disc of radius $r$, with the standard closed-form area $2\pi(\cosh r - 1)$. This diverges as $r \to \infty$, and these discs exhaust $S$ by completeness in the same way the $R_n$ did.

The two bounds contradict each other directly: the same quantity $\text{Area}(S)$ is at once $\leq 2\pi$ and unbounded above. No consistent choice is available, so the standing assumption — that a complete, regular, constant-negative-curvature $S$ exists at all — must fail. □

**Epistemic Notice**

Lemmas (a)–(d) above are reported in outline, following Hilbert's own short proofs of each rather than reproduced with full $\varepsilon$–$\delta$ care, and the geodesic-disc area formula $2\pi(\cosh r - 1)$ is quoted as the standard closed-form area of a hyperbolic disc, not re-derived here from chapter 10's own Poincaré-metric definition of $\mathbb{H}^2$. Both are exactly the kind of quoted-not-derived step this volume flags rather than hides elsewhere (e.g. section 10.4); what has changed from this chapter's earlier draft is that the argument's shape and every one of its steps now come directly from Hilbert's own 1901 text, not from a secondary paraphrase of unknown fidelity.

**Key Result**

*[The proof's logical skeleton, now read from Hilbert's own text] Appendix A's appendix A.9 record the access and reading behind the proof sketch above. Hilbert's own four lemmas about asymptotic curves, reported as (a)–(d) above, replace an earlier draft of this section that reported an eight-lemma exponential-map-and-Cartan–Hadamard structure from a modern secondary source (do Carmo/Spivak): building a global isometry from $\mathbb{H}^2$ onto $S$'s tangent space via the exponential map, then upgrading a local Tchebyshef-net parametrization to a global diffeomorphism in five further steps. That secondary account is not wrong exactly — it is a legitimate, more modern route to the same global-chart conclusion (a)–(d) establish — but it is a different argument from Hilbert's own, and this chapter now reports Hilbert's, consistent with the rest of this volume's preference for a primary source once one is in hand. One question this comparison surfaced, previously left open here: Hilbert's own paper attributes the quadrilateral area formula to Darboux, stated there as* $\text{Area} = (\text{angle sum}) - 2\pi$*, apparently the mirror image, in sign, of an earlier draft of proposition 11.6 in this chapter,* $\text{Area} = 2\pi - (\text{angle sum})$*. A first guess — that the two are reconciled by an interior/exterior angle relabelling — turned out on inspection not to be the mechanism: section 11.4's from-scratch derivation shows the discrepancy instead traced to this chapter's own earlier, uncredited choice of sine-Gordon sign convention when proposition 11.6 was first drafted from do Carmo's secondary account, before section 11.3's own concrete kink solution and verified sign were pinned down independently. With that derivation now anchored to this chapter's own verified* $\omega_{uv} = \sin\omega$ *and to explicit corner-angle bookkeeping rather than to either secondary source, proposition 11.6 has been restated as* $\text{Area} = (\text{angle sum}) - 2\pi$*, matching Hilbert/Darboux directly; this item is accordingly closed, not merely narrowed.*

**Todo — verify primary source**

Theorem 11.8's bibliographic citation is independently corroborated by five secondary sources consulted before the primary text itself was located (an Archive for History of Exact Sciences paper, a Perception journal bibliography, an arXiv bibliographic compilation, Wikipedia's own account, and an EPFL knowledge-graph entry), all agreeing on *Transactions of the American Mathematical Society* 2(1), 87–99 — now additionally confirmed directly against the article's own JSTOR-hosted scan (appendix A.9). That same earlier pass located two further items not previously on this chapter's list: a second, independent proof published the following year by E. Holmgren, "Sur les surfaces à courbure constante négative" (1902), and a far-reaching generalization by Nikolai Efimov, extending Hilbert's $C^2$ result to a broader regularity setting.
A later pass located both precisely, though neither has been read. Holmgren's paper is confirmed, from five independent secondary citations agreeing exactly, to be *Comptes Rendus de l'Académie des Sciences* (Paris), volume 134 (1902), pages 740–743; the journal's full pre-1923 run, including this volume, is digitized and openly accessible through the Bibliothèque nationale de France's Gallica platform (back catalogue 1835–1965), so access is a located, precisely scoped future step, not an open search. A serendipitous find from the same search: Édouard Goursat published an independent, admittedly more restricted, and by his own account considerably simpler proof of the same non-realizability result — "Sur les surfaces à courbure constante négative," *Bulletin de la Société Mathématique de France* 37 (1909), pages 51–58 — freely available in full via Numdam, the open French mathematical archive; this is a third proof route not previously on this chapter's list at all, and, unlike Holmgren's, sits at an access level (open, no paywall, no request needed) better than either Hilbert's or Holmgren's own papers required.
The Efimov citation this chapter originally carried here, as a 1975 *Vestnik MGU* paper, does not survive contact with chapter 12's own considerably more thorough bibliographic work: that chapter's section 12.2 traces Efimov's actual paper, cross-confirmed by two independent secondary sources' own reference lists, to N. V. Efimov, "Generation of singularities on surfaces of negative curvature" (Russian), *Matematicheskii Sbornik* 64 (1964), pages 286–320 — not 1975, and not *Vestnik MGU* — with an English translation confirmed to exist, at a precise volume and page range, in *American Mathematical Society Translations*, Series 2, volume 66 (1968), pages 154–190, though it remains unread behind a paywall. This chapter's own citation is accordingly retired in favour of chapter 12's: the 1975/*Vestnik MGU* attribution was itself never independently verified here, and chapter 12 both narrows and corrects it. Holmgren's paper remains open, precisely enough for a future pass to attempt.

## 11.6 A second theorem in the same paper: rigidity of the sphere

Hilbert's 1901 paper is in two parts, not one. The first, reported above, treats negative curvature. The second, headed (in the original) *Ueber Flächen von positiver constanter Krümmung*, answers the positive-curvature analogue of a different question: not whether $S^2$ can be built at all (chapter 2's own subject already answers that trivially, by exhibiting it), but whether it can be built in more than one way.

**Theorem 11.9** (Hilbert, 1901, after Liebmann, 1899)**.** *Let $\Sigma$ be a closed, regular, analytic surface of constant Gaussian curvature $+1$ in $\mathbb{R}^3$. Then $\Sigma$ is the unit sphere.*

*Proof sketch, following Hilbert's own argument.* Hilbert credits the closed case's reduction to H. Liebmann (*Göttinger Nachrichten*, 1899, p. 44) without reproving it, and proves himself a local statement that does the real work: on any singularity-free piece of an analytic surface of curvature $\equiv +1$, the larger of the two principal curvature radii $r_1 \geq r_2$ (with $r_1 r_2 \equiv 1$) attains its maximum only on the piece's boundary — unless the piece is (part of) the unit sphere, where $r_1 \equiv r_2 \equiv 1$ throughout and every point is an umbilic.

Suppose instead that $r_1$'s maximum, strictly greater than 1, is attained at an interior point $O$. Then $O$ is not an umbilic, so the two principal curvature directions foliate a neigh-

bourhood of $O$; use them as coordinate curves $(u, v)$ centred at $O$. Writing $\rho = \frac{1}{2}\log\frac{r_1+1}{r_1-1} > 0$, the Gauss equation for curvature-1 surfaces in principal coordinates gives the Liouville-type equation $\rho_{uu} + \rho_{vv} = e^{-2\rho} - e^{2\rho}$, quoted by Hilbert without further derivation (structurally the elliptic counterpart of proposition 11.5's hyperbolic sine-Gordon equation). Since $\rho$ decreases as $r_1$ increases, $\rho$ has a local *minimum* at $O$: its Taylor expansion there has no linear term, and its quadratic part $\alpha u^2 + 2\beta uv + \gamma v^2$ is positive semidefinite, forcing $\alpha, \gamma \geq 0$. Substituting the expansion into the Liouville equation at $(u, v) = (0, 0)$ gives $2(\alpha + \gamma) = e^{-2\rho(O)} - e^{2\rho(O)} < 0$, since $\rho(O) > 0$ makes $e^{-2\rho(O)} < e^{2\rho(O)}$ — directly contradicting $\alpha, \gamma \geq 0$. No interior maximum with $r_1 > 1$ is possible, which is the local statement. □

**Key Result**

*[Read directly from Hilbert's own text, not from any secondary source] The theorem and proof above are reported from appendix A's appendix A.9 reading of Hilbert's own 1901 paper. This material is new to the volume as of this pass and was not previously reported anywhere in it, including in section 11.7's own $g = 0$ bullet below, which asserted the sphere's rigidity in passing, without proof or citation, before this pass.*

**Todo — verify primary source**

Liebmann's own 1899 closed-case paper, cited but not reproved by Hilbert and not reproved here either, remains unread; so does Hilbert's own closing claim, stated without proof, that the local maximum-principle lemma just proved — granted for a closed surface with no boundary on which a maximum could hide — immediately forces $r_1 \equiv 1$ everywhere and hence the full rigidity conclusion. This is recorded as a new open item, of the same shape as section 11.4's and section 11.5's own quoted-not-derived steps.

A later pass narrowed, without closing, the Liebmann half of this. The paper's precise title is now known — not merely "Liebmann, Gött. Nachr. 1899, p. 44," as Hilbert's own footnote and every secondary source consulted here give it, but H. Liebmann, "Eine neue Eigenschaft der Kugel" ("A new property of the sphere"), *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1899, page 44 — confirmed against the journal's own table of contents as reproduced by the German Wikisource project, which links each entry directly to the Göttinger Digitalisierungszentrum's (GDZ) own freely accessible scan; the item is digitized and openly available with no paywall, login, or lending session, on the same public-domain footing appendix A.9 above established for Hilbert's own paper. A further, unconfirmed but plausible reconciliation located in the same pass: Liebmann's 1899 note is very likely a short announcement rather than the full proof, since standard bibliographies of the period cite an evidently fuller treatment by the same author immediately afterward — H. Liebmann, "Über die Verbiegung der geschlossenen Flächen positiver Krümmung," *Mathematische Annalen* 53 (1900), pages 81–112, described in multiple independent sources as Liebmann's Leipzig habilitation thesis of 1899 — following exactly the announcement-then-full-paper pattern chapter 12's own bibliographic work found for Efimov's 1963 *Doklady* note versus his 1964 *Matematicheskii Sbornik* paper. Neither the 1899 note's nor the 1900 paper's actual content has been read in this pass; both are recorded here precisely enough for a future pass to attempt, and Hilbert's own unproved reduction claim remains entirely untouched by this narrowing, since it is Hilbert's own logical gap, not a citation gap that locating Liebmann's paper would close.

## 11.7 What this adds to the volume's picture of curvature sign

**Corollary 11.10.** *For every $g \geq 2$, the closed hyperbolic surface $\Sigma_g$ of theorem 10.3 — compact, and therefore automatically complete — admits no isometric $C^2$ immersion into $\mathbb{R}^3$.*

This corollary is immediate from theorem 11.8, since compactness is a special case of completeness, but it is worth stating explicitly because it sharpens corollary 10.4's trichotomy

in a way that trichotomy's own intrinsic statement cannot see. Intrinsically — as far as theorem 2.1 alone can tell — the three cases $g = 0, 1, g \geq 2$ are on equal footing: each admits a constant-curvature metric, and in each case that metric obstructs a flat map by exactly the same mechanism. Extrinsically — asking whether the abstract metric can be realized as an actual complete surface sitting in ordinary space — the three cases split apart sharply:

- $g = 0$: the round sphere of chapter 2's own subject realizes $K \equiv +1/R^2$ completely, smoothly, and without incident, as the literal unit sphere in $\mathbb{R}^3$ — and, by theorem 11.9 above, is moreover the *only* closed regular analytic surface that does so: unlike $g \geq 2$'s outright non-realizability, the positive-curvature case is realizable but rigid, admitting no second, exotic embedding to sit alongside the standard one.
- $g = 1$: the flat torus realizes $K \equiv 0$ completely and smoothly, but — a fact this volume flags without proof, since it lies outside its own subject of curvature obstructions — not *isometrically as an ordinary smooth ($C^2$) embedded surface* in $\mathbb{R}^3$ in the everyday doughnut shape; the familiar embedded torus of revolution is not a flat metric (it has positive curvature on its outer equator and negative curvature on its inner one, a fact chapter 9's own Gauss–Bonnet computation already required to average to zero). A genuinely flat torus does admit an isometric $C^1$ embedding in $\mathbb{R}^3$, by a celebrated and highly non-classical construction (Borrelli, Jabrane, Lazarus, and Thibert, 2012, building on the Nash–Kuiper $C^1$ isometric embedding theorem), but no $C^2$ one, by an argument of the same convexity-theoretic flavor as theorem 11.8's own.
- $g \geq 2$: corollary 11.10 forbids even a $C^2$ isometric immersion, complete or otherwise once compactness is granted, into $\mathbb{R}^3$ at all.

**Todo — verify primary source**

The $g = 1$ bullet above is included for contrast and is not fully independently verified by this chapter to the standard the rest of this volume applies. One part of it has since been narrowed by a later pass (recorded in chapter 12, drafted after this chapter): N. Kuiper's own 1955 result — "On $C^1$-isometric imbeddings I, II," *Nederlandse Akademie van Wetenschappen, Proceedings, Series A*, volume 58 (*Indagationes Mathematicae* 17), pages 545–556 and 683–689 — is cited, by a paper this volume's Efimov chapter reads in full, as showing precisely that regularity $C^2$ (as opposed to $C^1$) is essential even in Hilbert's own theorem, which is a specific, citable source for claim (b)'s regularity-threshold half. That citation had not been independently obtained or read at the time this chapter was drafted; a later pass, recorded in chapter 13's own section 13.6, has since read Kuiper's Part I directly, in È. G. Poznyak's 1957 Russian translation rather than the English original (section 13.6.1 states what that licenses), confirming the regularity-threshold claim from the primary text at that standing. Claim (a) — that no smooth torus of revolution is isometric to the flat torus — remains reported from general knowledge only, though chapter 13's own lemma 13.9 and corollary 13.11 now supply a considerably more general from-scratch proof of the same fact (no compact $C^2$ surface of any genus is flat), and a further pass has obtained and read the precise content of the 2012 Borrelli–Jabrane–Lazarus–Thibert construction in full (section 13.7), rather than reporting it from general knowledge as this footnote originally did. This item is accordingly closed by chapter 13's later passes, not merely narrowed, consistent with section 11.4's own flagged citation gap above.

The sharper picture, then, is this. Chapter 10 showed that curvature sign is invisible to the specific obstruction theorem 2.1 states — $K > 0$ and $K < 0$ block a flat map identically. Theorem 11.8 shows that curvature sign is very much visible to a different, harder question this volume had not previously asked: whether the surface can exist, complete and unbounded, in ordinary space at all. The sphere can be built. The hyperbolic plane, in any of its constant-curvature-$(-1)$ incarnations, closed or open, cannot — not because no one

has found the right construction yet, but because theorem 11.8 proves none exists. This is the volume's first impossibility theorem whose obstruction is neither the pointwise curvature mismatch of theorem 2.1 nor the topological index argument of chapter 6, but a genuinely global area/angle bound (proposition 11.6) that only bites once a surface is asked to extend without limit — a third mechanism, alongside the volume's other two, for a cartographic-adjacent impossibility to arise.

**Todo — verify primary source**

No primary source for theorem 11.8 itself had been consulted at the time this chapter was first drafted: David Hilbert, "Über Flächen von constanter Gaussscher Krümmung," *Transactions of the American Mathematical Society*, volume 2, number 1 (1901), pages 87–99, was named only from the standard citation given in the secondary literature this chapter otherwise relied on. A later pass, recorded in appendix A's appendix A.9, located, fetched, and read the paper in full via JSTOR's Early Journal Content program (freely accessible, no paywall or lending session); section 11.5's proof sketch, section 11.5's keyresult, section 11.6's new section, and this synthesis's $g = 0$ bullet above have all since been rewritten or added to reflect that reading directly, rather than left pointing at a source not yet consulted. What the rewrite left open at that stage, precisely rather than generally: section 11.5's sign discrepancy between Hilbert/Darboux's own quadrilateral-area formula and proposition 11.6 as this chapter stated it. A later pass has since closed that specific item: section 11.4 re-derives proposition 11.6 from scratch from this chapter's own sine-Gordon equation and corner-angle bookkeeping, confirms the derivation both symbolically and by 40-digit numerical quadrature, and finds it forces the Hilbert/Darboux sign — proposition 11.6 has been restated accordingly. Proposition 11.6's own Hazzidakis-1880 citation gap (section 11.4) remains untouched by this, since Hilbert's paper turns out not to name Hazzidakis at all, and the derivation just closed is independent of, not a substitute for, locating Hazzidakis's own 1880 text. Section 11.6's own closing todobox still records that Liebmann's 1899 closed-case paper (now precisely titled and located, freely accessible via GDZ, but still unread) and Hilbert's own unproved reduction from the local lemma to the full closed case both remain open. Holmgren (1902), the further citation section 11.5's corroboration pass located, is now precisely located too (*Comptes Rendus* 134, pages 740–743, freely accessible via Gallica) but likewise unread; the chapter's earlier Efimov (1975, *Vestnik MGU*) citation has been retired in favour of chapter 12's own corrected and better-corroborated 1964 *Matematicheskii Sbornik* citation.

# Chapter 12

# Efimov's Theorem: When Constant Curvature Was the Easy Case

## 12.1 Cohn-Vossen's question

Chapter 11 closed with a clean trichotomy: the round sphere can be built in $\mathbb{R}^3$, the flat torus can be built (though not smoothly, past a regularity threshold chapter 11's own closing todobox left unresolved), and the hyperbolic plane cannot be built at all, at any regularity level *theorem* 11.8 demands. But *theorem* 11.8 proves something narrower than its informal statement in chapter 10 suggested: it forbids a complete surface of *constant* curvature $K \equiv -1$, not a complete surface whose curvature merely stays negative. The gap between those two hypotheses is not a technicality. A surface whose curvature is negative everywhere but drifts toward 0 far from some basepoint is a fundamentally more flexible object than one held rigidly at $K \equiv -1$, and nothing in *section* 11.5's asymptotic-net argument — built entirely on the sine-Gordon equation $\omega_{uv} = \sin\omega$, which is specific to constant $K = -1$ — says anything about it.

This is exactly the question Stefan Cohn-Vossen raised in 1936, a generation after Hilbert: does completeness plus curvature *bounded above by a negative constant* — $K \leq -c < 0$ everywhere, for some fixed $c > 0$, allowing $K$ to vary from point to point so long as it never approaches 0 — still forbid a $C^2$ isometric immersion into $\mathbb{R}^3$? Cohn-Vossen's own conjecture was that it does; the conjecture resisted proof for nearly three decades, and the eventual resolution, due to Nikolai Efimov, is the subject of this chapter.

**Key Result**

*[Primary source located and the conjecture's own statement confirmed] A further pass located a freely accessible full-text scan of Cohn-Vossen's own 1936 paper, hosted by the Russian mathematical portal Math-Net.Ru at* `https://www.mathnet.ru/rus/rm5945` *(OCR'd, in Russian; the scan itself is imperfect, with garbled diacritics and occasional dropped letters, but legible). The paper's own title, confirmed directly rather than taken on a secondary source's word, is "Изгибаемость поверхностей в целом"* (Bendability of surfaces in the large)*, and its overwhelming majority (§§1–6, roughly forty of its forty-four pages) is a survey of rigidity and uniqueness results for* convex *surfaces (Cauchy's theorem, Hilbert's uniqueness proof for the sphere, Blaschke's rigidity proof for ovaloids, Rembs's results on partial spheres and flat-strip surfaces) — not, as this chapter's own framing might suggest, centrally about negative curvature at all. The specific conjecture this chapter attributes to Cohn-Vossen appears only in the paper's final section (§7, "Некоторые гипотезы," "Some hypotheses"), as one among several speculative extensions Cohn-Vossen offers without proof. Located and translated directly from the primary text, the passage states (in this chapter's own translation, not a reproduction of the Russian): a hypothesis that has already occupied many mathematicians without result, that apparently no closed surface without singularities exists in $\mathbb{R}^3$ whose Gaussian curvature satisfies $K \leq -1$ at every point — offered explicitly as the negative-curvature analogue of Hilbert's uniqueness theorem for the sphere, in the same way Cohn-Vossen's own uniqueness conjecture for general convex ovaloids (this paper's actual main subject) is offered as the analogue of uniqueness for the sphere specifically. This confirms, from the primary source directly, both that the conjecture is genuinely Cohn-Vossen's own and that his own stated form used the specific bound $K \leq -1$ rather than a general $K \leq -c < 0$ — mathematically equivalent up to rescaling, since $c > 0$ can always be absorbed by scaling the ambient metric, but worth recording as the precise form actually written in 1936.*

**Epistemic Notice**

As with chapter 11, this chapter reports Efimov's theorem from secondary sources, not from Efimov's own paper, which is in Russian and has not been obtained or read in this pass; see the closing todobox. A first pass through this chapter (recorded in this volume's working log, not reproduced here) consulted two arXiv research papers only for their abstracts and introductions, and concluded no responsible proof sketch could be given. A further pass, reading both papers' full bodies rather than their abstracts, changed that conclusion: one of the two (Mendonça, building explicitly and closely on an earlier, considerably more detailed proof by Tilla Klotz Milnor) presents a complete, checkable, and reasonably short proof of theorem 12.1 itself, structured around the Gauss map rather than the sine-Gordon machinery of section 11.5. Section 12.4 below reports that argument's logical skeleton directly from Mendonça's own presentation, still a secondary source relative to Efimov's original 1964 paper, but a considerably more load-bearing one than a citation-only account. Section 12.5 below remain this chapter's own independent symbolic contributions, and section 12.1 is this chapter's own first direct reading of a primary source anywhere in this chapter.

## 12.2 Efimov's theorem

**Theorem 12.1** (Efimov, 1963)**.** *Let $M$ be a connected surface and $\varphi\colon M \to \mathbb{R}^3$ a $C^2$ isometric immersion such that $M$, with its induced metric, is complete. Then the Gaussian curvature $K$ of $M$ is* ***not*** *bounded above by a negative constant: there is no $c > 0$ with $K \leq -c < 0$ at every point of $M$.*

*Remark* 12.2. Theorem 11.8 is the special case of theorem 12.1 in which $K$ is not merely bounded above by a negative constant but held *exactly* at a single constant value $K \equiv -1$ everywhere; theorem 12.1 is correspondingly the strictly stronger statement, resolving Cohn-Vossen's 1936 conjecture and confirming that Hilbert's obstruction was never really about constancy of curvature, only about curvature staying uniformly negative. Theorem 12.1 says

nothing at all, by contrast, about a surface whose curvature is negative everywhere but not bounded away from 0 — and, as section 12.5 below shows concretely, such surfaces exist, are complete, and are entirely unremarkable.

**Todo — verify primary source**

A first pass gave this item's citation as Wikipedia's own reference, N. V. Efimov, "The impossibility in Euclidean 3-space of a complete regular surface with a negative upper bound of the Gaussian curvature," *Doklady Akademii Nauk SSSR*, volume 150, pages 1206–1209 (1963). A further pass, reading the full bodies (not just abstracts) of the two arXiv papers already cited above, found both papers' own reference lists cite Efimov's theorem to a different, more specific source: N. V. Efimov, "Generation of singularities on surfaces of negative curvature" (Russian), *Matematicheskii Sbornik*, volume 64 (1964), pages 286–320 — and both papers give this exact citation independently (Gálvez–Martínez–Teruel's reference [9], Mendonça's reference [Ef]), which this chapter treats as meaningfully corroborating, since the two papers' own reference lists were compiled independently of each other and of the Wikipedia citation found in the first pass. The most likely reconciliation, not independently confirmed, is that the 1963 *Doklady* item is a short announcement and the 1964 *Matematicheskii Sbornik* paper is Efimov's own full treatment with proofs — a common publication pattern in Soviet-era mathematics — but this chapter has not read either paper and cannot confirm the relationship directly. This chapter now cites the *Matematicheskii Sbornik* 1964 paper as primary, on the strength of its two independent corroborations, while keeping the 1963 *Doklady* citation on record as the likely announcement. Neither paper has been obtained or read in the original Russian.

A further pass (in this pass) found a genuinely new lead rather than a further corroboration of the existing citation: a bibliography in an arXiv preprint on pseudo-spherical surfaces lists an English-language translation of the 1964 paper itself, not merely a secondary description of it — N. V. Efimov, "The appearance of singularities on a surface of negative curvature," *Matematicheskii Sbornik* 64(106), no. 2 (1964), pages 286–320, English translation in *American Mathematical Society Translations*, Series 2, volume 66, pages 154–190. A later pass confirmed this precisely, directly from the publisher rather than from the single preprint citation: the AMS's own bookstore listing for *American Mathematical Society Translations, Series 2, Volume 66* ("Thirteen Papers on Group Theory, Algebraic Geometry and Algebraic Topology," 1968, 272 pages) gives its full table of contents, and the seventh of the thirteen papers is exactly N. V. Efimov, "Appearance of singularities on surfaces of negative curvature." The volume's own umbrella title is misleading on its face — group theory, algebraic geometry, and algebraic topology, with no mention of differential geometry — since these AMS Translations volumes are themselves miscellaneous collections of contemporaneous Soviet journal papers assembled by subject-adjacency rather than a single unified theme, and Efimov's paper sits among the thirteen despite not matching the umbrella title at all. The citation is accordingly no longer a single, uncorroborated lead but a publisher-confirmed fact: an English translation of Efimov's 1964 paper exists, in print, at this exact volume and page range. It remains unread: the AMS sells the volume as a $265 eBook with no free excerpt located, and no other free copy of specifically pages 154–190 was found in this pass. The item is accordingly narrowed from "does this translation exist" (now answered: yes) to "is it accessible without payment" (still open).

## 12.3 Why constant curvature was the easy case

Section 11.5's proof of theorem 11.8 leaned on one algebraic convenience that theorem 12.1's more general hypothesis destroys outright: constant curvature $K \equiv -1$ makes the Gauss equation in asymptotic coordinates collapse to the single, exactly solvable sine-Gordon equation $\omega_{uv} = \sin\omega$, and Hazzidakis's formula (proposition 11.6) then converts an area bound directly into an angle-sum bound with no further curvature-dependent correction term. Neither convenience survives when $K$ is allowed to vary, subject only to $K \le -c < 0$: the Gauss equation becomes $\omega_{uv} = -K(u,v)\sin\omega$ with $K$ now an unknown function rather than a fixed constant, and Hazzidakis's clean area-equals-angle-defect identity no longer applies verbatim,

since it was derived specifically for the constant-curvature case. Efimov's actual route around this, per section 12.4 below, abandons the asymptotic-net/Hazzidakis machinery entirely and argues instead through the Gauss map — a strategy section 11.5 did not need, since constant curvature let Hilbert stay entirely within the surface's own asymptotic coordinates.

## 12.4 Efimov's argument, via the Gauss map

**Epistemic Notice**

This section reports the proof's logical skeleton directly from Sérgio Mendonça's 2014 paper (section 12.8 below), which states explicitly that it follows the argument as written in Tilla Klotz Milnor's 1972 *Advances in Mathematics* paper — itself a detailed secondary presentation of Efimov's original argument, not Efimov's own text. This chapter has read Mendonça's short, self-contained proof of theorem 12.1 itself (as opposed to his own paper's main extension, Theorem A, which is more general and considerably longer) in full, and reports it here at the level of its main steps; the technical lemmas underlying the one load-bearing claim this sketch cannot avoid — that the Gauss map extends to an injective map on the whole surface — are not reproduced, for the same reason section 11.5's own disclaimerbox gave for Hilbert's theorem: they are the proof's genuine technical heart, established through a long sequence of convex-hull constructions on the sphere that this chapter has read but does not consider itself able to summarize responsibly in a few sentences.

Suppose, for contradiction, that $\varphi\colon M \to \mathbb{R}^3$ is a $C^2$ isometric immersion of a connected surface $M$, complete in the induced metric $\alpha$, with Gaussian curvature $K \le -\kappa < 0$ everywhere. Let $N\colon M \to S^2$ be the Gauss map of $\varphi$, and let $\beta$ be the pullback to $M$ of the round metric $\nu$ on $S^2$ under $N$ — concretely, $\beta$ is the third fundamental form of $\varphi$, exactly the object section 11.5's own proof of theorem 11.8 also introduced, though section 12.4 here puts it to different use.

**Proposition 12.3** (Area transforms by $|K|$)**.** $A_\beta(M) = \int_M \bigl|\det(dN)\bigr|\, dA_\alpha = \int_M |K|\, dA_\alpha$.

This is immediate from the definition of pullback area and the classical fact that the Gauss map's Jacobian determinant is exactly the Gaussian curvature (the same identity chapter 6's own treatment of the Gauss map uses elsewhere in this volume, there for the index-counting argument rather than an area computation).

Since $K < 0$ everywhere, $N$ is a local diffeomorphism (section 12.4's hypothesis rules out the parabolic points where $\det(dN) = 0$ that would obstruct this), so $M$ is simply connected: any complete, simply connected surface of non-positive curvature is a standard fact of comparison geometry (reported here, not re-derived, following chapter 10's own use of the same Cartan–Hadamard theorem for $\mathbb{H}^2$), and it carries a further standard consequence this proof needs directly: *a complete, simply connected surface of non-positive Gaussian curvature has infinite area.* Combined with $K \le -\kappa < 0$, this gives $\int_M |K|\, dA_\alpha \ge \kappa \cdot A_\alpha(M) = \infty$, so by proposition 12.3, $A_\beta(M) = \infty$ as well.

The contradiction is that $A_\beta(M)$ can also be shown finite — and this is exactly the step section 12.4's own disclaimerbox flagged as unreproduced here. The argument shows, through the long induction on convex neighborhoods already flagged above, that $N$ is in fact *injective* on all of $M$ (not merely a local diffeomorphism), and a local isometry onto its image inside $(S^2, \nu)$; since $N$ is injective and $S^2$ has finite total area $4\pi$, this forces $A_\beta(M) = A_\nu\bigl(N(M)\bigr) \le A_\nu(S^2) = 4\pi < \infty$. Infinite by the curvature bound, finite by injectivity of the Gauss map: no such $\varphi$ exists.

*Remark* 12.4. The shape of this argument is worth setting directly against section 11.5's: both proofs derive a contradiction from an area bounded two ways, one growing without limit and one capped by a fixed constant ($2\pi$ there, $4\pi$ here). Section 11.5 capped the area of an

ever-larger *coordinate* quadrilateral via Hazzidakis's formula, intrinsic to the surface's own asymptotic net; section 12.4 instead caps the area of the *entire surface*, pulled back through the Gauss map, by the fixed total area of the unit sphere. The two proofs are not the same argument in different notation — section 11.5's Hazzidakis bound has no direct analogue here, and is exactly the ingredient section 12.3 above noted does not survive variable curvature — but both belong to the same family this volume's synthesis (sections 11.7 and 12.9) keeps returning to: a growing quantity trapped underneath a fixed ceiling.

**Todo — verify primary source**

The single load-bearing gap in this section is the injectivity of $N$ on all of $M$, which this chapter has read (in Mendonça's Lemmas 4.1–4.7, itself explicitly following Klotz Milnor's considerably longer 1972 treatment) but does not reproduce. At the level of its own stated strategy, the argument builds injectivity outward from a small geodesic ball by an induction that extends a convex, $N$-injective neighborhood along successive short geodesic arcs, using the surface's completeness and boundary-free-ness (formalized as the non-existence of "concave" points on the metric completion of $M$ under $\beta$) to rule out the extension ever failing partway; the induction is closed by a separate argument (attributed to Myers's theorem in the surface-with-boundary case Mendonça's own paper treats, not needed in quite that form for the boundaryless case sketched here) bounding how far the construction must reach before wrapping around $S^2$. This is reported as a roadmap of the argument's shape, consistent with section 11.5's own disclaimerbox's standard for what this volume calls a "proof sketch," not as an independently checked derivation. This chapter has also not independently verified the standard fact this proof leans on — that complete, simply connected, non-positively-curved surfaces have infinite area — beyond noting that both chapter 10 and Mendonça's own paper cite it as standard.

## 12.5 The hypothesis is sharp: the catenoid

Theorem 12.1's hypothesis is exactly "$K \le -c < 0$," not the visually similar but strictly weaker "$K < 0$ everywhere." The gap between the two is not a matter of proof-technique convenience; it is load-bearing, and a single classical, elementary surface makes this concrete without needing anything beyond chapter 2's own first- and second-fundamental-form machinery.

**Definition 12.5** (Catenoid)**.** For $c > 0$, the *catenoid* of waist radius $c$ is the surface of revolution

$$X(u, v) = \big(c\cosh(u/c)\cos v,\ c\cosh(u/c)\sin v,\ u\big), \qquad u \in \mathbb{R},\ v \in [0, 2\pi).$$

**Proposition 12.6.** *The catenoid has Gaussian curvature*

$$K(u) = \frac{-1}{c^2\cosh^4(u/c)},$$

*strictly negative at every point, is a complete $C^\infty$ surface, and satisfies $K(u) \to 0$ as $u \to \pm\infty$: its curvature is **not** bounded away from zero on either end.*

**Key Result**

*[Verified by direct symbolic computation] Proposition 12.6 is confirmed by* `scripts/efimov_catenoid_curvature.py`*, which computes the catenoid's first and second fundamental forms directly from the parametrization above ($E = \cosh^2(u/c)$, $F = 0$, $G = c^2\cosh^2(u/c)$; $e = -1/c$, $f = 0$, $g = c$), forms $K = (eg - f^2)/(EG - F^2)$ symbolically, and confirms the result simplifies* exactly *to $-1/(c^2\cosh^4(u/c))$ — not merely numerically close to it, since sympy's symbolic simplification reduces the difference between the computed expression and this closed form to exactly 0. The script goes on to confirm three further, independent things numerically and symbolically rather than by assertion: (i) $K(u) < 0$ at nine sample points spanning $u \in [-10, 10]$; (ii) $\lim_{u\to\pm\infty} K(u) = 0$ exactly, via sympy's own limit computation; and (iii) that for every $\varepsilon > 0$ there exist points with $|K(u)| < \varepsilon$ (worked explicitly for $\varepsilon = 10^{-3}$, giving the threshold $|u| > 2.4121\ldots$), which is precisely the formal statement that $K$ is not bounded away from 0 on either end. A fourth, separate check confirms completeness along meridians specifically: the induced-metric arclength element along a fixed-$v$ curve is $\sqrt{E} = \cosh(u/c)$ (at $c = 1$), and the script computes the definite integral $\int_0^L \cosh(u)\,du = \sinh(L)$ symbolically, confirming it diverges as $L \to \infty$ — so a meridian of the catenoid has infinite length in both directions, exactly the completeness property theorem 12.1's hypothesis requires and section 11.2's pseudosphere conspicuously lacked.*

*Remark* 12.7. The catenoid is therefore a complete, $C^\infty$, everywhere-negatively-curved surface immersed (in fact embedded) in $\mathbb{R}^3$ — and this is entirely consistent with theorem 12.1, not a counterexample to it, precisely because its curvature fails the theorem's "bounded away from zero" hypothesis at both ends. The catenoid's two ends flatten out (its curvature decays toward the flat plane's $K = 0$ as $u \to \pm\infty$, visibly: the surface widens and its meridian straightens), and it is exactly this flattening — ruled out by theorem 12.1's hypothesis, permitted here — that lets completeness and negative curvature coexist where theorem 11.8's constant-curvature hypothesis, and theorem 12.1's bounded-away-from-zero hypothesis, both forbid it.

**Todo — verify primary source**

Section 12.5's verification is this chapter's own independent symbolic contribution, not drawn from a secondary source for the computation itself. A first pass through this section did not know whether the catenoid was the standard example cited alongside theorem 12.1 for this sharpness point; a further, full-body reading of Mendonça's paper (section 12.8 below) answers this directly and reveals the standard example is a *different* one: Mendonça's own Remark 1 states plainly that "the helicoid shows that the condition $K \leq -\kappa < 0$ may not be replaced by the condition $K < 0$," citing no further source, as though the fact were common knowledge in the field (which, given the helicoid's status as one of the two or three most classical non-trivial minimal surfaces, it plausibly is). This chapter's catenoid illustration is therefore its own choice, not a reproduction of the literature's own standard example, and the two are not obviously interchangeable: the helicoid's curvature decays with distance from its axis rather than along two flattening ends, a structurally different route to the same non-boundedness-away-from-zero conclusion.

**Key Result**

*[Verified by direct symbolic computation, and corrected] A further pass ran the same symbolic method section 12.5 used for the catenoid on the helicoid, via* `scripts/efimov_helicoid_curvature.py`*, parametrizing* $X(r,\theta) = (r\cos\theta, r\sin\theta, a\theta)$ *for pitch parameter* $a > 0$ *and radius* $r \geq 0$*. The first and second fundamental forms come out to* $E = 1$, $F = 0$, $G = a^2+r^2$; $e = 0$, $f = -a/\sqrt{a^2+r^2}$, $g = 0$, *giving* $K = (eg-f^2)/(EG-F^2)$*. This did* not *confirm the formula this chapter had previously cited without derivation,* $K = -1/(a^2+r^2)^2$*: the symbolic computation instead gives* $K = -a^2/(a^2+r^2)^2$ *exactly, a genuine correction (the two formulas coincide only at* $a = 1$*, which is presumably how the error entered — a transcription silently specialized to unit pitch). The corrected formula is confirmed exactly by sympy's simplification of the difference to* 0*, and the same further checks section 12.5 ran for the catenoid are run here too:* $K(r) < 0$ *at eight sample points from* $r = 0$ *to* $r = 50$ *(*$a = 1$*);* $\lim_{r\to\infty} K(r) = 0$ *exactly;* $|K(r)|$ *is strictly decreasing in* $r$ *(confirmed via* $dK/dr = 4a^2r/(a^2+r^2)^3 > 0$ *for* $r > 0$*, so* $K$ *itself increases toward* 0*); for* $\varepsilon = 10^{-3}$, $|K(r)| < \varepsilon$ *once* $r > 5.5338\ldots$*, exhibiting the same non-boundedness-away-from-zero failure of theorem 12.1's hypothesis the catenoid exhibits, here by radial decay around a single axis rather than by two flattening ends; and the radial arclength element is* $\sqrt{E} = 1$ *identically, so* $\int_0^L 1\,dr = L \to \infty$*, confirming completeness along radial lines specifically, in the same limited sense section 12.5's meridian check did for the catenoid.*

**Todo — verify primary source**

Section 12.5's correction was found by this chapter's own independent symbolic derivation, not by consulting a second secondary source for the formula — so while the derivation itself is checked exactly (sympy reduces the discrepancy and the confirmation to exact zero, not a numerical near-match), no external source has yet been consulted to confirm that $-a^2/(a^2+r^2)^2$, rather than $-1/(a^2+r^2)^2$, is indeed the form standard references use; the two are mathematically equivalent only up to the meaning assigned to $a$, and it remains possible the original citation used a different, also-valid parametrization convention (e.g. absorbing $a$ elsewhere) rather than being a bare error. This chapter treats its own from-scratch derivation as authoritative for the specific parametrization stated above, but has not chased down why the earlier, uncorrected form appeared in this chapter's own text in the first place, nor checked it against Mendonça's paper or any other named source, since Mendonça's own Remark 1 (section 12.5's main text) states the qualitative fact without giving a formula at all.

## 12.6 A more dramatic failure of sharpness: Nadirashvili's surface

The catenoid's curvature merely *decays* toward zero at two isolated ends; it does not come arbitrarily close to 0 on some complicated, unbounded set threading through the whole surface. A considerably more dramatic illustration of theorem 12.1's sharpness is due to Nikolai Nadirashvili: N. Nadirashvili, "Hadamard's and Calabi–Yau's conjectures on negatively curved and minimal surfaces," *Inventiones Mathematicae* 126, number 3 (1996), pages 457–465. A further pass, prompted by dissatisfaction with the two-sentence, single-Wikipedia-source account this section originally gave, located and read (in full) a 2024 survey by Antonio Alarcón that treats Nadirashvili's theorem as its own historical starting point, and this section is substantially rewritten on that basis.

Nadirashvili's theorem settles two older questions at once, not one. The first, due to Jacques Hadamard in 1898, asked directly whether a complete bounded immersed surface with negative Gaussian curvature could exist in $\mathbb{R}^3$ at all. The second is a special case ($n = 3$) of a 1965 conjecture of Eugenio Calabi — promoted independently by Shing-Tung Yau, who restated it as an open problem in 1982 in terms this literature calls the *Calabi–Yau problem* — that no complete bounded *minimal* hypersurface exists in $\mathbb{R}^n$ at all, curvature sign aside. Nadirashvili's 1996 paper answers both together, in the negative, for $n = 3$:

**Theorem 12.8** (Nadirashvili, 1996)**.** *There is a complete conformal minimal immersion of the open unit disc $\mathbb{D} \subset \mathbb{C}$ into $\mathbb{R}^3$ whose image lies in the unit ball. Moreover, such an immersion exists with everywhere-negative Gaussian curvature.*

**Key Result**

*[The construction's shape, reported from a detailed secondary source] Alarcón's survey describes Nadirashvili's construction as an inductive refinement of a compact minimal disc, using three classical tools in combination: the Weierstrass representation of minimal surfaces, the López–Ros deformation (a method for deforming a conformal minimal immersion while preserving one coordinate function), and Runge's approximation theorem for holomorphic functions applied on an increasingly intricate "labyrinth" of compact subsets of the disc. The survey attributes the construction's key new idea to an elementary but carefully exploited consequence of the Pythagorean theorem: a boundary deformation that pushes each boundary point of a bounded piece of surface outward by a distance on the order of $s$, in a direction roughly orthogonal to that point's own position vector, increases the surface's* intrinsic *distance from a fixed interior point by an amount on the order of $s$ — but increases the* extrinsic *radius of the ball containing the surface only on the order of $s^2$, since the new radius is $\sqrt{r^2+s^2}$ rather than $r+s$. Iterating this imbalance across an infinite recursive sequence of such deformations, with the perturbation size shrinking fast enough at each step to converge to a limit immersion, produces a sequence of intrinsic boundary distances diverging to infinity while the extrinsic containing radius converges to a finite limit — exactly theorem 12.8's completeness-inside-a-bounded-ball conclusion.*

*Remark* 12.9 (Self-intersection is not incidental). Nadirashvili's immersion is not embedded, and the survey confirms this is not a gap a cleverer construction could close: Tobias Colding and William Minicozzi (*Annals of Mathematics* (2), volume 167, number 1 (2008), pages 211–243) proved that a complete *embedded* minimal surface of finite topology in $\mathbb{R}^3$ must be proper — and a proper immersion into $\mathbb{R}^3$ is automatically unbounded, since properness means preimages of compact sets are compact, which forces the surface to escape every bounded region. A complete bounded *embedded* minimal surface of finite topology is therefore impossible outright, not merely unconstructed; William Meeks, Joaquín Pérez, and Antonio Ros later extended the same properness conclusion to surfaces of finite genus with countably many ends (*Duke Mathematical Journal*, volume 170, number 13 (2021), pages 2891–2956). Whether a complete bounded *embedded* minimal surface of infinite genus or with uncountably many ends could exist is, per the survey's own account, still an open question as of its 2024 writing — this chapter has not sought a more recent update on that specific point.

**Todo — verify primary source**

This section is substantially improved from its earlier two-sentence, single-secondary-source state, but Nadirashvili's own 1996 paper has still not been obtained or read directly; section 12.6's account of the construction, and theorem 12.8's precise statement, are both taken from Alarcón's 2024 survey (itself citing Nadirashvili's paper, and a separate expository note by Collin and Rosenberg, *Bulletin des Sciences Mathématiques* 123, number 7 (1999), pages 563–575, neither independently consulted here). No computation in this section has been independently verified by this chapter in the manner section 12.5's catenoid computation was; the Weierstrass-representation/López-Ros machinery underlying the actual construction is considerably more involved than the pseudosphere or catenoid parametrizations this volume has verified directly elsewhere, and reproducing even a toy instance of it is left as a lead for a future pass, not attempted here.

## 12.7 What is possible in higher dimensions

One further data point sharpens theorem 12.1's specifically three-dimensional character. Chapter 11 already noted, in its own closing synthesis, that the flat torus (curvature iden-

tically 0, not negative) admits a $C^1$ but not $C^2$ isometric embedding in $\mathbb{R}^3$ — regularity, not dimension, was the loophole there. For the hyperbolic plane specifically, the loophole is dimension: *raising* the ambient dimension past 3 restores what theorem 11.8 forbids in $\mathbb{R}^3$. Two constructions, both predating Efimov's own theorem and both concerning $\mathbb{H}^2$ itself (constant curvature, theorem 11.8's setting, not theorem 12.1's more general one), are now cross-confirmed with precise citations in both papers' own reference lists: Danilo Blanuša's 1955 smooth complete isometric *embedding* of $\mathbb{H}^2$ in $\mathbb{R}^6$ ("Über die Einbettung hyperbolischer Räume in euklidische Räume," *Monatshefte für Mathematik* 59 (1955), pages 217–229), and Èrnst Rozendorn's 1960 smooth complete isometric *immersion* of $\mathbb{H}^2$ in $\mathbb{R}^5$ ("A realization of the metric $ds^2 = du^2 + f^2(u)$ in a five-dimensional Euclidean space" (Russian, Armenian summary), *Akademiya Nauk Armyanskoi SSR, Doklady* 30 (1960), pages 197–199). A further data point, found only in Mendonća's own reference list and not previously on this chapter's radar, extends theorem 12.1 itself to higher dimensions directly, rather than exhibiting an escape from it: B. Smyth and F. Xavier, "Efimov's theorem in dimension greater than two," *Inventiones Mathematicae* 90 (1987), pages 443–450, proving (per Mendonća's own one-sentence description, not independently read) that no codimension-one complete isometric immersion $f\colon M \to \mathbb{R}^n$ exists with Ricci curvature $\mathrm{Ric}_M \leq -\kappa < 0$, when $n = 3$ or when $n \geq 4$ and the sectional curvatures of $M$ do not take every real value.

**Todo — verify primary source**

None of the three items in this section — Blanuša (1955), Rozendorn (1960), Smyth–Xavier (1987) — has been obtained, read, or independently verified in this pass; all three citations are now precise (title, journal, volume, page range) and, for the first two, cross-confirmed by both of this chapter's two main secondary sources' independent reference lists, but the underlying constructions and proofs remain entirely unread. Whether an analogous escape — a complete $C^2$ isometric immersion, for a variable-but-bounded-away-from-zero-negative curvature satisfying theorem 12.1's hypothesis rather than $\mathbb{H}^2$'s constant one, in some ambient dimension between 3 and the threshold Smyth–Xavier's own generalization would allow — is known to exist is not addressed by either source as consulted, and is left here as an open question this chapter does not attempt to answer even at the level of a citation.

## 12.8 Later extensions, briefly

Chapter 11's own closing todobox flagged, from a bibliographic pass consulted before this chapter was drafted, a 1975 paper by Efimov himself (in Russian, *Vestnik MGU*) extending his own result to a broader regularity setting than $C^2$; this pass did not separately re-locate or cross-check that citation, and it remains exactly as open as chapter 11 left it. Two further extensions, both now read in full rather than at abstract level, are recorded here for their actual content.

Gálvez, Martínez, and Teruel (2016) prove that a complete $C^2$-immersed surface with $K \leq -\kappa < 0$ holding only *outside a compact set* (rather than everywhere) still has finite total curvature, finite area, is properly immersed, and has every end asymptotic to a half-line in $\mathbb{R}^3$ — a genuinely more general conclusion than theorem 12.1's own outright non-existence, since a compact region is now permitted to violate the curvature bound however it likes. The paper's technical route generalizes the same Gauss-map machinery section 12.4 above sketches (its own Generalized Lemma 1 is explicitly the planar-domain result underlying Efimov's original proof, restated to accommodate a compact-boundary case) rather than replacing it. The same paper takes up, separately and only partially, John Milnor's own 1972 conjecture that a complete surface with principal curvatures satisfying $k_1^2 + k_2^2 \geq \text{const} > 0$ must either change the sign of $K$ or be flat — proving, as a partial case, that non-positive curvature together with one principal curvature bounded away from zero forces the surface

to be a generalized cylinder — and extends both the Efimov-type and Milnor-type results to immersions in hyperbolic 3-space $\mathbb{H}^3$ and the round 3-sphere $S^3$, via an abstract theory of "Codazzi pairs" (a metric and a second, Codazzi-compatible quadratic form on an abstract surface, generalizing the first and second fundamental forms of an actual immersion).

Mendonça (2014) proves, as his own paper's Theorem A, that no complete $C^2$ immersion of a non-compact surface $M$ with compact boundary into $\mathbb{R}^3$ can satisfy both $\int_M |K| = +\infty$ and $K \leq -\kappa < 0$ — a hypothesis genuinely different in character from Gálvez–Martínez–Teruel's (infinite total curvature in place of a compact exceptional region), from which theorem 12.1 itself is recovered as the boundaryless special case, and which the paper's own Corollary B specializes further to complete Hadamard (open, simply connected, non-positively curved) surfaces. Section 12.4 above is this chapter's own direct report of the proof Mendonça gives for the classical, boundaryless case of theorem 12.1 on the way to this more general Theorem A; the more general boundary case adds a further, longer layer of geodesic-convexity lemmas (his own Lemmas 4.1–4.7) this chapter has read but does not separately summarize, since section 12.4's own todobox already reports the boundaryless argument's corresponding step at the same level of detail.

**Todo — verify primary source**

Both papers are now read in full rather than at abstract level, closing the specific gap this todobox previously flagged. What remains open: neither paper's own proofs have been independently re-derived or symbolically checked by this chapter (unlike section 12.5's catenoid computation); the Milnor conjecture Gálvez–Martínez–Teruel address is stated here exactly as their own introduction states it, without this chapter separately verifying John Milnor's original 1972 formulation against a source distinct from their citation of it; and the Codazzi-pairs machinery underlying both papers' extensions to $\mathbb{H}^3$ and $S^3$ is reported only in outline, not worked through.

## 12.9 What this adds to the volume's picture

Chapter 11 closed by naming three mechanisms this volume's impossibility theorems rest on: the pointwise curvature mismatch of theorem 2.1, the topological index argument of chapter 6, and Hilbert's own global area/angle bound. Theorem 12.1 does not add a fourth mechanism — its proof, so far as section 12.3 above can responsibly describe it, is a variable-coefficient descendant of Hilbert's own area/angle argument, not a structurally new one — but it sharpens what that third mechanism actually needs. Theorem 11.8 might have left the impression that the obstruction to building the hyperbolic plane in $\mathbb{R}^3$ was somehow tied to the perfect rigidity of *constant* curvature, a special algebraic convenience rather than a geometric necessity. Theorem 12.1 rules that impression out directly: the obstruction survives intact under the much weaker, and geometrically far more natural, hypothesis that curvature merely stay uniformly negative. Section 12.5's catenoid then draws the boundary precisely, in a way section 12.3's proof-difficulty discussion alone could not: it is not negativity of curvature that Efimov's mechanism needs, but negativity held uniformly away from the flat case $K = 0$ that the sphere and the flat torus, in their own ways, already occupy at the trichotomy's other two points. A surface is free to build itself in $\mathbb{R}^3$, completely and smoothly, exactly to the extent its curvature is willing to approach 0 somewhere — at every point, for the flat torus; nowhere, but only in the limit, at the catenoid's two ends; and, per section 12.6's unverified but striking report, on some far more complicated unbounded set for Nadirashvili's construction. Uniform separation from $K = 0$, not the sign of $K$ alone, is what theorems 11.8 and 12.1 together identify as the true extrinsic obstruction.

**Todo — verify primary source**

No primary source for theorem 12.1 itself has been consulted in this pass. N. V. Efimov, "Generation of singularities on surfaces of negative curvature" (Russian), *Matematicheskii Sbornik*, volume 64 (1964), pages 286–320, is named as the paper's precise bibliographic identification per section 12.2's todobox — corrected and strengthened, in this pass, from a single-source *Doklady* 1963 citation to a citation cross-confirmed by two independent papers' own reference lists — but the paper itself, in Russian, in a Soviet-era journal not obviously covered by the pre-1923 American-public-domain window appendices A.4 and A.5 exploited elsewhere in this volume, has not been located, fetched, or read; nor has the 1936 Cohn-Vossen paper conjecturing the result (*Uspekhi Matematicheskikh Nauk* 1 (1936), pages 33–76, also newly precise in this pass), nor Tilla Klotz Milnor's 1972 *Advances in Mathematics* paper (volume 8, pages 474–543) that section 12.4 above leans on at one remove via Mendonça's own presentation of it. This pass did close, rather than merely narrow, three items from this list: section 12.8's two extension papers, previously flagged as read only at abstract level, have now been read in full; section 12.6's Nadirashvili account has been substantially rewritten from a two-sentence, single-Wikipedia-source placeholder to a report grounded in a full reading of a 2024 survey; and section 12.5's symbolic check of the helicoid's curvature formula, previously left as a suggested future step, both closes that step and corrects the formula it was meant to verify. This pass also narrowed, without fully closing, three further items. Section 12.2's todobox now records the AMS Translations citation as publisher-confirmed rather than merely found in one preprint's bibliography: the translation demonstrably exists, at the exact volume and page range stated, though it remains unread behind a paywall. Section 12.1 converts Cohn-Vossen's 1936 paper from wholly unread to located, freely accessible, and read at a targeted (not comprehensive) level: the specific conjecture-bearing passage has been found, translated, and confirmed directly from the primary source, though the paper's other six sections — the bulk of it, on convex-surface rigidity rather than negative curvature — remain unread beyond a skim for context. What remains open, gathered in one place: Efimov's own 1963 *Doklady* note (still entirely unread) and the 1964 *Sbornik* paper (Russian original unread; an English translation is now confirmed to exist in print, at AMS Translations Series 2, vol. 66, pp. 154–190, but sits behind a $265 paywall with no free copy located); the bulk of Cohn-Vossen's 1936 paper beyond the one located passage; Klotz Milnor's 1972 paper; Nadirashvili's own 1996 paper (only the 2024 survey's report of it has been read); the Blanuša, Rozendorn, and Smyth–Xavier constructions (section 12.7); and independent confirmation (beyond this chapter's own unaided translation) that the located Cohn-Vossen passage is standardly read the way this chapter reads it, since no second source has cross-checked this chapter's own translation of the Russian.

# Chapter 13

# The Nash–Kuiper Theorem: Regularity as the Loophole

## 13.1 An open thread from two earlier chapters

Chapter 11 closed its treatment of the flat torus with a claim offered "for contrast" and explicitly flagged as not independently verified to this volume's usual standard: that the flat torus $\mathbb{R}^2/\mathbb{Z}^2$ admits an isometric $C^1$ embedding in $\mathbb{R}^3$, by a construction built on the *Nash–Kuiper $C^1$ isometric embedding theorem*, even though theorem 11.8's own machinery — and, more directly, the elementary fact this chapter proves in section 13.8 below — forbids any $C^2$ realization of the same metric. Chapter 12 returned to the same claim once, narrowing it slightly (attaching a precise citation, N. Kuiper's 1955 two-part paper, to the regularity-threshold half) without resolving it. Both chapters left the same two things undone: no chapter of this volume states the Nash–Kuiper theorem itself, and no chapter explains *why* weakening $C^2$ to $C^1$ defeats an obstruction as robust as theorems 11.8 and 12.1's, which chapter 12's own closing synthesis (section 12.9) had just finished describing as a fully general statement about uniform separation from $K = 0$. This chapter takes up both tasks, and adds a third: an independent proof, not previously given anywhere in this volume, of the specific claim — claim (a) of chapter 11's $g = 1$ discussion — that no smooth torus of revolution, or indeed any compact $C^2$ surface of any genus, can be isometric to the flat torus at all.

**Epistemic Notice**

This chapter's account has been substantially upgraded from an earlier pass, in which none of its three primary sources had been obtained. Nash's 1954 paper and the first of Kuiper's two 1955 papers have now been read in full, in a specific, precisely identified form: not the English originals, but È. G. Poznyak's Russian translations, published together in *Matematika* 1, no. 2 (1957), pages 3–16 (Nash) and 17–28 (Kuiper, Part I), freely available from the Steklov Mathematical Institute's Math-Net.Ru portal. Sections 13.5 and 13.6 report their content at this standing — a contemporary, edited translation of the complete argument, not a paraphrase filtered through any secondary exposition — and section 13.6.1 states plainly what this standing does and does not license. A further pass has since obtained and read Borrelli, Jabrane, Lazarus, and Thibert's 2012 paper in full as well, in its own published English original; section 13.7 now reports its convex integration mechanism, its Corrugation Theorem, and its own reported numerical figures directly, rather than from secondary description. Only Kuiper's second 1955 paper (pages 683–689, on non-compact manifolds) remains unread among this chapter's four named primary sources. Section 13.4's toy corrugation model and section 13.8's farthest-point lemma remain this chapter's own independent mathematical content, unaffected by the above: the first is verified directly by `scripts/nash_kuiper_corrugation.py`, the second is proved in full from first principles below.

## 13.2 Short maps, and the shape of the theorem

**Definition 13.1** (Short map)**.** Let $(M, g)$ be a Riemannian manifold and $f\colon M \to \mathbb{R}^n$ a $C^1$ immersion. Write $f^*e$ for the pullback to $M$ of the Euclidean metric $e$ on $\mathbb{R}^n$ under $f$. Then $f$ is *short* if $f^*e \le g$ as quadratic forms at every point — that is, $|df_p(v)| \le \sqrt{g_p(v,v)}$ for every tangent vector $v$ — with *strict* inequality on a dense set. Equivalently, $f$ never stretches any tangent direction beyond what $g$ prescribes, and typically stretches it strictly less.

**Theorem 13.2** (Nash, 1954; Kuiper, 1955)**.** *Let $(M^m, g)$ be a smooth Riemannian manifold and $f_0\colon M \to \mathbb{R}^n$ a short $C^\infty$ immersion (embedding) with $n \ge m+1$. Then for every $\varepsilon > 0$ there exists a $C^1$ isometric immersion (embedding, if $f_0$ was one) $f\colon M \to \mathbb{R}^n$ — meaning $f^*e = g$ exactly, not merely $f^*e \le g$ — with $\sup_{p\in M} |f(p) - f_0(p)| < \varepsilon$.*

*Remark* 13.3. The bound $n \ge m+1$ is the striking part of the statement. Nash's own original 1954 result required $n \ge m+2$ for embeddings (one dimension of ambient slack beyond the codimension-one case); Kuiper's 1955 refinement is specifically the sharpening to $n = m+1$ — the *same* codimension-one setting *theorems* 11.8*and* 12.1 both live in, surfaces ($m = 2$) in ordinary space ($n = 3$). This is exactly why theorem 13.2 bears on this volume's own preceding two chapters rather than sitting one dimension removed from their concerns: Kuiper's contribution was precisely to bring the flexibility theorem down into the ambient dimension where chapters 11 and 12's rigidity theorems operate.

**Key Result**

*Two ordinary embedded surfaces already available elsewhere in this volume instantiate definition 13.1's hypothesis directly, without any new construction. Chapter 2's stereographic or Lambert projections, run in reverse as a map from a flat disc into $\mathbb{R}^3$ landing back on the sphere, are not short in general; but a simple uniform scaling is: if $(M, g)$ is any compact Riemannian 2-manifold and $f_0$ is* any *smooth embedding into $\mathbb{R}^3$, then $f_0$ composed with a sufficiently small uniform contraction $\lambda f_0$, $0 < \lambda < 1$ small enough, is short in the sense of definition 13.1, since contracting Euclidean distances by $\lambda$ while $g$ itself is held fixed eventually shrinks every pulled-back length strictly below what $g$ prescribes. Theorem 13.2 therefore applies to* any *compact surface with any smooth Riemannian metric $g$ and any smooth embedding into $\mathbb{R}^3$ at all: shrink first, then corrugate. This is the sense in which the theorem's hypothesis is easy to satisfy — the entire content of theorem 13.2 is what happens next.*

## 13.3 Why this does not contradict Hilbert or Efimov

Chapters 11 and 12 both derive their obstructions from second-order data: section 11.5's proof of theorem 11.8 works through the asymptotic net and Hazzidakis's area formula, both built from the Gaussian curvature $K$, itself a second-derivative quantity of the embedding by chapter 2's own definitions; section 12.4's proof of theorem 12.1 works through the Gauss map $N\colon M \to S^2$, whose differential — the object whose Jacobian determinant *is* $K$ (proposition 12.3) — is again second-order in the embedding $\varphi$. A $C^1$ immersion has a well-defined first derivative (hence a well-defined induced metric $f^*e$, hence a well-defined notion of "isometric") but no well-defined second derivative at all: curvature, in the classical pointwise sense chapter 2's Theorema Egregium is a statement about, simply is not defined for a merely $C^1$ map. Neither theorem 11.8's nor theorem 12.1's proof has any object to apply to.

*Remark* 13.4 (Flexibility is not an accident of one clever construction). Theorem 13.2 is not, on the level of secondary description this chapter relies on, an isolated trick specific to embeddings into Euclidean space. It is among the founding examples of what M. Gromov later organized into the general theory of the *h-principle* ("h" for homotopy): a broad class

of geometric problems, phrased as differential relations, for which existence of *some* solution reduces — once the relation is, in a suitable technical sense, "open" and "ample" — to existence of a solution to the far weaker, purely algebraic-topological relaxation obtained by forgetting that the derivatives in question need to come from an actual map at all (here: the existence of a short map, which is easy, standing in for the existence of an isometric one, which is what is actually wanted). This chapter does not attempt to state Gromov's convex integration machinery in general, consistent with the standing this volume gives to results reported at citation level rather than independently re-derived; it records the connection because it is the conceptual reason theorem 13.2's conclusion — a dense abundance of wild, nowhere-twice-differentiable isometric embeddings, coexisting with theorems 11.8 and 12.1's complete absence of any $C^2$ one — is not a curiosity restricted to this one theorem, but a recurring pattern wherever a rigidity theorem's proof genuinely needs second-order information and the corresponding existence question is asked one derivative down.

## 13.4 The mechanism in miniature: corrugating a line segment

Nash's actual proof builds $f$ as a limit of a sequence of "stage" maps, each obtained from the last by adding a further, higher-frequency *corrugation* — a wave-like perturbation, transverse to the previous stage's image, that increases arc length locally without moving the map far in the sup norm. The full construction, iterated across a decomposition of $g - f_0^{*}e$ into a finite sum of manageable pieces (Nash's own device), is well beyond what this chapter can responsibly summarize from secondary knowledge alone. But the one-dimensional case of the same idea — corrugating a single curve to lengthen it while staying $C^0$-close to a shorter one — is elementary, exactly computable, and independently verifiable; it is not a simplification borrowed from any source but this chapter's own worked instance of the mechanism's shape.

**Proposition 13.5** (A corrugated curve of prescribed length)**.** *Fix $A > 0$ and a positive integer $N$, and define $\gamma_N\colon [0,1] \to \mathbb{R}^2$ by*

$$\gamma_N(t) = \Big(t,\ \tfrac{A}{2\pi N}\sin(2\pi N t)\Big).$$

*Then $\gamma_N$ is a smooth embedded curve with:*

*(i) exact arc length $L(A) = \dfrac{2}{\pi}\sqrt{1+A^2}\,E\!\left(\dfrac{A^2}{1+A^2}\right)$, where $E$ is the complete elliptic integral of the second kind, **independent of** $N$;*

*(ii) sup-norm distance $\sup_t |\gamma_N(t) - (t,0)| = \dfrac{A}{2\pi N} \to 0$ as $N \to \infty$;*

*(iii) derivative $\gamma_N'(0) = (1, A)$ for every $N$, never approaching the straight segment's derivative $(1,0)$.*

*Proof.* $\gamma_N'(t) = \big(1,\, A\cos(2\pi N t)\big)$, so the speed is $|\gamma_N'(t)| = \sqrt{1 + A^2\cos^2(2\pi N t)}$, and $L(A,N) = \int_0^1 \sqrt{1 + A^2\cos^2(2\pi N t)}\,dt$. The integrand is periodic in $t$ with period $1/N$, and $[0,1]$ is exactly $N$ periods, so the substitution $s = Nt$ gives $L(A,N) = \int_0^1 \sqrt{1 + A^2\cos^2(2\pi s)}\,ds$, manifestly independent of $N$; writing $\cos^2\theta = 1 - \sin^2\theta$ and factoring out $\sqrt{1+A^2}$ converts this into 4 times the standard integral defining $E$ over a quarter period, giving the stated closed form. (ii) is immediate from $|\sin(2\pi N t)| \le 1$. (iii) is $\gamma_N'(0) = (1, A\cos 0) = (1, A)$ for every $N$, with no $N$-dependence at all. □

**Key Result**

*[Independently verified] Proposition 13.5 is confirmed by* `scripts/nash_kuiper_corrugation.py` *at several levels. First, the closed-form length is cross-checked against direct numerical quadrature of the $N$-dependent integrand for $A \in \{0.3, 0.8, 1.5, 3.0\}$ and $N \in \{1, 5, 37, 200, 1000\}$, agreeing to better than $10^{-8}$ in every one of the twenty combinations tested (quadrature subdivision budget scaled with $N$ to resolve all $N$ oscillation periods; an initial run with a fixed, $N$-independent subdivision budget failed exactly the high-$N$ cases in the way this predicts, a useful negative check confirming the test itself is sensitive to the effect it is meant to catch). Second, the sup-norm distance is confirmed numerically equal to $A/(2\pi N)$ to within sampling resolution, for $N$ ranging over four orders of magnitude (1 to $10^4$). Third, $\gamma_N'(0) = A$ exactly, for every tested $N$ up to $10^4$, confirming the derivative gap neither shrinks nor depends on $N$. Fourth, the root-finding direction — given a target length $L_{\text{target}} \in \{1.05, 1.3, 1.8, 2.5\}$, solve for $A$ via* `scipy.optimize.brentq` *against the closed form, then confirm the resulting curve achieves exactly that length at several $N$ — succeeds to within $10^{-12}$ in every case.*

*Remark* 13.6 (What this miniature does and does not show). Proposition 13.5 exhibits, in the simplest setting where it is exactly computable, the trade this chapter's title names: for *any* target length $L \geq 1$, there is a smooth curve — for every $N$, not merely in some limit — lying within $A/(2\pi N)$ of the straight unit segment, whose length is exactly $L$. Sending $N \to \infty$ makes the curve converge uniformly ($C^0$) to the straight segment while its length stays fixed at the target value the whole time; the price, per (iii), is that the curve's direction at $t = 0$ never approaches the straight segment's own direction, for any $N$ — so no subsequence converges in $C^1$, only in $C^0$. This is the one-dimensional shadow of theorem 13.2's own trade: the limit map is $C^1$ (each finite-$N$ $\gamma_N$ here is in fact $C^\infty$, and a genuine limit curve realizing the target length exactly, obtained by an appropriate diagonal argument across increasing $N$ and correspondingly finer corrugation, would be $C^1$ but generically no smoother), $C^0$-close to the short map it corrugates, but built by a mechanism — adding bounded-amplitude, unboundedly-fast oscillation — that makes no attempt to converge in $C^1$ and has no reason to be $C^2$ anywhere.

## 13.5 Nash's own construction: twisting deformations

Nash's paper — *Matematika* 1:2 (1957), 3–16, translating *Annals of Mathematics* (2) 60 (1954), 383–396 — proves theorem 13.2 for $n \geq m + 2$ by a construction that matches this chapter's section 13.4 in spirit but not in every detail this chapter had guessed. The process is an infinite sequence of *stages*, each stage subdivided into finitely many *twists* (Poznyak's Russian renders Nash's term as "скручивания"; this chapter's earlier draft, working only from secondary accounts, had already guessed "corrugation" as a reasonable English gloss, which the primary text now confirms is the right idea under a different word). Each stage is assigned, via a smooth partition of unity $\{\varphi_p\}$ subordinate to a locally finite cover of $M$ by coordinate neighborhoods $N_p$, a share of the total metric deficit $g - h$ (where $h$ is the current induced metric); within each $N_p$, that share is further decomposed into finitely many twists by the following device, stated and proved as Nash's own Lemma 1.

**Key Result**

*[Nash's Lemma 1, reported from the primary text] Let $\sigma$ be a positive-definite symmetric $m \times m$ matrix (representing, pointwise, the metric correction to be supplied). Nash covers the open convex cone of positive-definite symmetric $m \times m$ matrices by open simplicial neighborhoods, each lying entirely inside the cone and meeting only finitely many others (a star-finite cover, with the maximum multiplicity $W$ depending only on $m = \frac{1}{2}n'(n'+1)$ where $n' = m$), so that $\sigma$, as an interior point of one such simplex, is a positive-coefficient combination of the positive-definite rank-one matrices at that simplex's vertices; a $C^\infty$ partition of unity subordinate to this cover, applied to $\sigma$ itself as it varies smoothly over $N_p$, converts this into a smooth decomposition*

$$\sigma = \sum_\nu a_\nu(x)\, \psi_\nu \otimes \psi_\nu, \qquad a_\nu \geq 0 \text{ smooth}, \quad \psi_\nu \text{ linear in local coordinates},$$

*with only finitely many terms nonzero near any point of $N_p$.*

This is a genuinely topological device — a simplicial covering of the cone of positive-definite matrices, exactly analogous to the covering of $M$ itself by coordinate patches — and not, as this chapter's earlier pass had impressionistically described it, a bare "decomposition into manageable pieces" left unspecified. Each term $a_\nu \psi_\nu \otimes \psi_\nu$ in section 13.5's decomposition is realized by one twist: writing $\lambda$ for a large parameter chosen per twist and $\xi, \eta$ for a pair of smooth unit orthonormal vector fields normal to the current immersion (whose existence, for $N \geq m+2$, Nash derives from Steenrod's cross-section theorems for sphere bundles, cited to *Topology of Fibre Bundles*), the twisted immersion is

$$z^\alpha \;=\; z_0^\alpha \;+\; \frac{\xi^\alpha}{\lambda}\cos(\lambda\psi_\nu) \;+\; \frac{\eta^\alpha}{\lambda}\sin(\lambda\psi_\nu),$$

matching, up to notation, this chapter's own section 13.4 miniature (the toy curve $\gamma_N$ is exactly the $m = 1$, single-normal-field shadow of this formula with $\eta$ absent). Nash computes the induced metric change directly: terms linear in the normal fields vanish by normality ($\xi \cdot dz_0 = \eta \cdot dz_0 = 0$), cross terms in $\xi, \eta$ vanish by their mutual orthogonality, and the surviving $\cos^2 + \sin^2 = 1$ terms combine to give exactly $a_\nu\, d\psi_\nu \otimes d\psi_\nu$ up to an error of order $O(1/\lambda)$ — the same big-$O$ bookkeeping this chapter's proposition 13.5 verified numerically in miniature, now confirmed as the actual shape of Nash's own error term, not merely an analogous toy phenomenon.

*Remark* 13.7 (Why Nash needs *two* normal fields). The twist formula uses $\xi$ and $\eta$ together, at the same time, as the two legs of a genuine rotation in the normal 2-plane they span: the point moves along a circle of radius $1/\lambda$ in that plane as $\lambda\psi_\nu$ varies. This is the structural reason theorem 13.2's original 1954 proof needs codimension at least two ($N \geq m+2$): the deformation is intrinsically two-dimensional in the normal directions it consumes. Section 13.6 below reports Kuiper's replacement of this two-field twist by a one-field mechanism, which is exactly the "why" this chapter's earlier disclaimerbox had left as an open item.

The remainder of Nash's paper — convergence of the stage process (a geometric-series argument bounding the accumulated first-derivative error by a factor of $\sim 5/6$ per stage, driving $C^1$ but not $C^2$ convergence, exactly matching remark 13.6's toy-model description), the embedding case (preventing self-intersection by a limit-set argument controlling each twist's displacement against a fixed exhaustion by compact sets), and the construction of an initial short embedding for open manifolds (an explicit $5(n+2)$-dimensional construction using a locally finite cover, class-separated to avoid overlap, with monotonically shrinking scale constants $\varepsilon_p$) — is reported here as read and confirmed, not paraphrased from any secondary source. The four numbered theorems Nash states at the paper's close match, term for term, what theorem 13.2 above already asserted from secondary knowledge: closed

$n$-manifolds admit $C^1$ isometric embeddings in $E^{2n}$; general $n$-manifolds admit $C^1$ isometric immersions in $E^{2n}$ and embeddings in $E^{2n+1}$; the codimension-two threshold for compact manifolds is sharp in Nash's own construction (he explicitly flags, in closing, that he expects $k = 1$ to suffice by a harder argument — the exact conjecture section 13.6 reports Kuiper as resolving). Nash's own worked illustration, credited in his closing paragraph to Milnor, is the flat $n$-torus itself: C. Tompkins's 1939 *Duke Mathematical Journal* result forbids a $C^2$ (in fact Tompkins's result is stated for real-analytic) isometric embedding of the flat torus in dimension below $2n$, while an ordinary $C^\infty$ embedding into $E^{n+1}$ is easy to write down directly — so a short rescaling of that embedding already gives, by theorem 13.2, a $C^1$ isometric embedding in $E^{n+2}$, one full dimension below Tompkins's smooth bound. This is, verbatim, the flat-torus phenomenon chapter 11 first flagged and this chapter's own section 13.8 independently reproves below in the $n = 2$ case — now traced to Nash's own paper as its origin, not merely a modern gloss on it.

The chapter's original section 13.4 description of Nash's decomposition as "a finite sum of manageable pieces (Nash's own device)" was vague in a way now worth flagging explicitly: it correctly anticipated that *some* finite decomposition exists, but had no access to what made it work. Section 13.5's simplicial covering of the positive-definite matrix cone is the actual mechanism, and it is a more interesting fact than "some decomposition exists" — it is itself a small application of the same partition-of-unity technology used at the level of the manifold $M$, one dimensional level up, applied to the space of possible metric corrections rather than to $M$ itself.

## 13.6 Kuiper's refinement: stretching corrugations replace twists

Kuiper's Part I — *Matematika* 1:2 (1957), 17–28, translating *Proceedings Koninklijke Nederlandse Akademie van Wetenschappen* A 58 (1955), 545–556 — opens by stating explicitly what it is answering: Nash's own closing conjecture that the codimension-two hypothesis could be weakened to codimension one, for compact manifolds with boundary. Kuiper proves exactly this, and states his method as deliberately independent of Nash's proof rather than a modification of it, because the finer convergence analysis his one-normal-field mechanism requires would not obviously follow by simply adapting Nash's two-normal-field error estimates.

**Key Result**

*[Kuiper's mechanism, reported from the primary text] Where Nash's twist rotates the point in the 2-plane spanned by* two *normal fields $\xi, \eta$ (remark 13.7), Kuiper's "stretching corrugation" (Poznyak's Russian: "**гофрировка с растяжением**") uses only* one*: writing $\xi$ for a single smooth unit normal field and $h = |\partial z_0/\partial x^1|$, the deformation is*

$$z^\beta \;=\; z_0^\beta \;-\; \frac{\frac{1}{8}\alpha^2 \sin 2\lambda x^1}{\lambda}\, h\xi^\beta \;+\; \alpha\Big(1 - \tfrac{1}{8}\alpha^2 \cos 2\lambda x^1\Big) \cos\Big[\lambda x^1 - \tfrac{1}{8}\alpha^2 \sin 2\lambda x^1\Big] \frac{\xi^\beta}{\lambda} \;+\; h\eta^\beta,$$

*built from the one-dimensional toy Kuiper works out first (his own miniature, reported by this chapter's section 13.4 independently and by coincidence rather than by borrowing): replacing a straight segment $y = 0$ by the plain sine wave $(x, a \sin x)$ concentrates the induced-metric change unevenly, large at inflection points and vanishing at crests and troughs, so Kuiper adds a second, compensating oscillation at twice the frequency (the $\alpha^2 \sin 2\lambda x^1$ term) to spread the metric change evenly along the curve — a refinement with no counterpart in Nash's twist, and the reason Kuiper calls his deformation a* stretching corrugation *rather than a twist: it uses the available normal direction more efficiently, at the cost of a harder convergence argument, precisely trading the second normal field for extra analytic care.*

*Remark* 13.8 (This closes the "why" this chapter's earlier pass had left open). Because section 13.6's deformation displaces the point using only one normal field $\xi$ at a time (the second field $\eta$ appearing only passively, carried along unchanged, in Kuiper's formula above), Kuiper's construction needs the ambient space to supply only one normal direction beyond the tangent space at each step, not two simultaneously as Nash's twist does. This is, now confirmed against Kuiper's own stated method rather than reconstructed from secondary paraphrase, exactly the mechanism this chapter's disclaimerbox had flagged as unverified: the reduction from $N \geq m+2$ to $N \geq m+1$ is not a sharper counting argument layered on top of Nash's construction, but a different deformation, one dimension leaner by design, substituting a harder one-dimensional convergence estimate (section 13.6's own toy calculation, carried through with explicit error bounds $O(\lambda^{-1})$ tracked via an approximation parameter Kuiper calls $\varepsilon_t$, shown to decrease geometrically stage over stage in his §9) for the room Nash's extra dimension had given him for free.

Kuiper's paper closes with the same structural results as Nash's — convergence of the deformed immersions to a $C^1$ limit (his own geometric-decrease argument for the error sequence $\varepsilon_t$, independent of Nash's), and a proof that the limit is free of self-intersections (via a covering-neighborhood argument bounding the ratio of true to Euclidean distance away from the diagonal of $M \times M$) — explicitly presented, per Kuiper's own introduction, as self-contained and not dependent on reading Nash's proof first. Kuiper's own headline example, stated in his introduction before any of the technical machinery, is the same flat torus: "the flat $n$-torus ... can be $C^1$-isometrically embedded in $E^{n+1}$" — one further dimension below Nash's own $E^{n+2}$ instance of the identical example, and, for $n = 2$, exactly $\mathbb{R}^2/\mathbb{Z}^2 \hookrightarrow \mathbb{R}^3$, the precise claim chapter 11 had originally flagged as unverified and this chapter exists to resolve.

**Todo — verify primary source**

Kuiper's second 1955 paper (pages 683–689, treating non-compact manifolds, per Nash's and Kuiper's own citation apparatus) has not been obtained or read; section 13.6's account is drawn entirely from Part I, whose own closing lines state that the non-compact extension is the subject of the sequel. Nothing in this chapter depends on Part II's specific content, since theorem 13.2 as stated only needs the compact-with-boundary case Part I already covers, but the item is left open rather than silently dropped.

### 13.6.1 What reading the translations, rather than the originals, licenses

Both papers read for this section are È. G. Poznyak's 1957 Russian translations, not Nash's or Kuiper's own English text, and the *Matematika* volume is explicit that Poznyak and an editor annotated the translation with footnotes of their own (one, attached to Nash's introduction, remarks that Nash's paper is harder to follow than Kuiper's and that the sharper constant $\frac{3}{5}$ could easily replace Nash's own $\frac{3}{4}$ at one point; another, attached to Kuiper's paper, suggests a specific sufficiently-large choice of the stage-index threshold $t_0$). These are visibly editorial interventions, distinguishable in the text from Nash's and Kuiper's own words, and this chapter has reported only the latter — the mathematical content of the constructions, statements, and proofs — not the editors' asides, as primary-source content. The translation is contemporary (1957, within three years of Nash's publication and two of Kuiper's), produced for a Soviet mathematical audience specifically because the results were recognized immediately as important, and edited by a named mathematician (Poznyak went on to publish his own work on isometric embeddings, cited elsewhere in the literature this volume has consulted); nothing in this chapter's reading suggests the translation is unfaithful to the mathematical content. But it remains a translation, not the archival original, and this chapter records that distinction rather than silently upgrading "read a faithful translation" to

"read the primary source" without qualification — the same standard of honesty this volume's appendix has applied throughout to sources consulted only in part or only secondarily.

## 13.7 The flat torus, corrugated explicitly

Chapter 11's own $g = 1$ discussion named the specific construction this section now reports from a direct reading, not a secondary paraphrase: Vincent Borrelli, Saïd Jabrane, Francis Lazarus, and Boris Thibert, "Flat tori in three dimensional space and convex integration," *Proceedings of the National Academy of Sciences* 109, number 19 (2012), pages 7218–7223 — located and read in full at a mirror of the publisher's own PDF (PNAS deposits its Direct Submissions in full-text open repositories; this pass's copy carries the paper's original pagination, author affiliations, and editorial handling line, "Edited by Yakov Eliashberg, Stanford University, Stanford, CA, and approved March 9, 2012," confirming it as the published version rather than a preprint draft).

**Key Result**

*[The construction's actual shape, not merely its existence] The paper's contribution is not citing theorem 13.2 abstractly for the flat torus — that existence was already known — but converting Gromov's convex integration machinery into an explicit, terminating algorithm and rendering its output. The mechanism differs from both sections 13.5 and 13.6's twist and stretching-corrugation constructions in one structural respect worth recording precisely: where Nash and Kuiper corrugate along a locally finite, essentially unbounded family of coordinate patches and directions, Borrelli, Jabrane, Lazarus, and Thibert observe that the isometric default of a well-chosen initial short embedding $f_{\text{init}}$ (an ordinary torus of revolution, shrunk) decomposes, at* every *point of the flat torus simultaneously, as a nonnegative combination of exactly three fixed rank-one quadratic forms $\ell_j \otimes \ell_j$, $j = 1, 2, 3$, built from three constant vector fields $U^{(1)} = e_1$, $U^{(2)} = \frac{1}{5}(e_1 + 2e_2)$, $U^{(3)} = \frac{1}{5}(e_1 - 2e_2)$ spanning the plane at angles chosen so their associated parallelograms tile $\mathbb{R}^2/\mathbb{Z}^2$. This lets the entire infinite corrugation process run as a single repeating block of exactly three one-dimensional-in-effect corrugations per stage, each canceling one of the three coefficients $\rho_{k,1}, \rho_{k,2}, \rho_{k,3}$ in turn, rather than requiring the patch-by-patch partition-of-unity bookkeeping section 13.5 needs. The one-dimensional corrugation formula itself (their Eq. 3, built from a barycentric condition matching a target speed $r$ to the average of a loop of unit-speed directions) is, up to notation, the same $O(1/N)$-controlled oscillation this chapter's own section 13.4 toy model and proposition 13.5's verification already illustrate in miniature — confirmed now against the actual published mechanism, not merely analogous to it.*

**Key Result**

*[The Corrugation Theorem: an explicit, if slowly convergent, fractal Gauss map] The paper's own Theorem 2 (their "Corrugation Theorem") expresses the frame rotation induced by a single corrugation step as a product of two orthogonal matrices, one ($L_{k,j}$) an $O(1/N_{k,j})$-accurate rotation by the oscillating corrugation angle $\theta_{k,j}$, the other ($R_{k,j}$) an $O(\varepsilon_{k,j})$-accurate rotation by the fixed angle between successive corrugation directions. Because the limiting Gauss map $n_\infty$ is then an infinite product of such near-rotation matrices, with amplitudes, frequencies, and directions varying slowly but never repeating exactly, the paper explicitly likens the structure to a Riesz product $\prod_k [1 + \alpha_k \cos(2\pi N_k x)]$, and states — without a full proof, flagged in their own text as likely but not fully established — that the Gauss map's graph plausibly has Hausdorff dimension strictly greater than two, even though the embedded torus itself, being a continuously differentiable injective image of a genuine 2-manifold, has Hausdorff dimension exactly two. This is the precise, technical content behind the qualitative "fractal normal vector" description this chapter's earlier secondary-sourced pass had already reported correctly in spirit; it is now traceable to a specific numbered theorem and a specific analogy (Riesz products) the paper itself draws, rather than to unspecified secondary paraphrase.*

**Key Result**

*[Concrete numbers: the corrugation counts explode, and the visualization is a lower-order approximation] The paper reports specific corrugation-frequency values for its own numerical implementation ($\gamma = 0.1$): $N_{1,1} = 611$, $N_{1,2} = 69{,}311$, $N_{1,3} = 20{,}914{,}595$, $N_{2,1} = 6{,}572{,}411{,}478$ — growing roughly two orders of magnitude per step, a rate the paper notes would require a grid of about $4.3 \times 10^{19}$ nodes for even ten samples per period by the fourth stage, well beyond feasible computation. The rendered image (their Fig. 3, not reproduced here) is accordingly* not *the full $C^1$ limit but a finite truncation after the first four corrugations (frequencies reduced by trial to 12, 80, 500, and 9,000 respectively, well below the theorem's sufficient-condition values, since those conditions are sufficient but not necessary); a direct length comparison the paper reports — meridians, parallels, and diagonals on the flat torus compared against their images under this truncated map — differs by at most 10.2% from the true flat-metric lengths, against an 80% deviation for the naive short torus-of-revolution map $f_{\text{init}}$ used as the starting point. This is a genuinely quantitative fact this chapter's earlier pass had no access to at all: not merely that the corrugated image is closer to isometric than the naive one, but by how much, at a specific, named stage of a specific, reproducible construction.*

## 13.8 Closing claim (a): no compact $C^2$ surface is flat

Chapter 11's $g = 1$ discussion left one further claim — that no smooth torus of revolution is isometric to the flat torus — "reported from general knowledge only." This section closes it, with a short, self-contained, and considerably more general proof: not merely tori of revolution, but *every* compact $C^2$ surface embedded in $\mathbb{R}^3$, regardless of genus, has curvature that cannot be identically zero.

**Lemma 13.9** (Farthest point has positive curvature)**.** *Let $\Sigma \subset \mathbb{R}^3$ be a compact, connected, $C^2$ surface without boundary. Then there is a point $p \in \Sigma$ at which the Gaussian curvature satisfies $K(p) > 0$.*

*Proof.* Fix any point $O \in \mathbb{R}^3$ and let $p \in \Sigma$ be a point maximizing the distance $|Op|$ over the compact set $\Sigma$ (such a $p$ exists since $\Sigma$ is compact and $|O\cdot|$ is continuous); write $r = |Op| > 0$. Since $\Sigma$ lies entirely within the closed ball $\bar{B}(O, r)$ and touches its boundary sphere $S(O, r)$ at $p$, the two surfaces are tangent at $p$: they share the same tangent plane, hence (up to a global sign convention on the normal, fixed by choosing the outward normal of the ball at $p$ for both) the same unit normal $\nu(p)$. For any unit tangent vector $v$ at $p$, compare the two surfaces' normal curvatures in direction $v$: because $\Sigma \subset \bar{B}(O, r)$ and both surfaces pass through $p$ tangentially, $\Sigma$ cannot curve away from the ball's interior faster than the sphere does in any direction — formally, the second fundamental form of $\Sigma$ at $p$ dominates that of $S(O, r)$ at $p$ (both computed with respect to the shared inward normal): $\mathrm{II}_\Sigma(v, v) \geq \mathrm{II}_{S(O,r)}(v, v) = \frac{1}{r}$ for every unit $v$. In particular every principal curvature of $\Sigma$ at $p$ is at least $1/r > 0$, so both principal curvatures are strictly positive there, and $K(p) = \kappa_1(p)\kappa_2(p) \geq 1/r^2 > 0$. □

*Remark* 13.10. This is a standard lemma (see, e.g., do Carmo, *Differential Geometry of Curves and Surfaces*, 1976, §3-3), reported here not as a new result but reproved from scratch, in the manner chapter 6's own from-scratch treatment of the Poincaré–Hopf index sum proceeded, rather than cited without derivation. The proof needs nothing beyond compactness and $C^2$-ness of $\Sigma$: not orientability, not genus, not any curvature hypothesis at all — it is purely an artifact of $\Sigma$ being closed and bounded in $\mathbb{R}^3$.

**Corollary 13.11.** *No compact $C^2$ surface embedded in $\mathbb{R}^3$, of any genus, has Gaussian curvature identically zero. In particular, the flat torus $\mathbb{R}^2/\mathbb{Z}^2$ — whose curvature is identically* 0 *by construction — admits no $C^2$ isometric embedding in $\mathbb{R}^3$: not the familiar torus of revolution, and not any other $C^2$ embedded genus-*1 *surface either.*

*Proof.* Immediate from lemma 13.9: a compact $C^2$ surface with $K \equiv 0$ would have no point of positive curvature, contradicting the lemma. Since an isometric embedding of the flat torus would, by definition, be a compact $C^2$ surface with $K \equiv 0$ everywhere (curvature is intrinsic, by theorem 2.1, so it is preserved exactly by any isometric embedding), no such embedding exists. □

*Remark* 13.12 (This sharpens, and generalizes, what chapter 11 had left informal). Chapter 11's own $g = 1$ bullet was stated only for the familiar embedded torus of revolution, and only as a flagged, unverified claim. Corollary 13.11 proves the considerably stronger statement that *no* compact $C^2$ embedded surface of *any* genus can have $K \equiv 0$ — a genus-0 sphere cannot either, nor could a genus-2 double torus, though for those two cases chapter 9's own Gauss–Bonnet computation ($\int_\Sigma K\, dA = 2\pi\chi(\Sigma)$, giving $4\pi \neq 0$ for the sphere and $-4\pi \neq 0$ for the double torus) already forbade $K \equiv 0$ by a global-integral argument, without needing lemma 13.9's pointwise one. The genus-1 case is exactly the one Gauss–Bonnet's integral *cannot* rule out on its own — $\chi = 0$ for the torus is consistent with $K \equiv 0$ at the level of the integral identity alone, since $\int_\Sigma 0\, dA = 0 = 2\pi \cdot 0$ — which is precisely why chapter 11 could only flag the $g = 1$ case as unverified rather than derive it from machinery already in the volume: chapter 9's tool is silent exactly where it is needed, and lemma 13.9's pointwise argument is what fills the gap.

## 13.9 What this adds to the volume's picture

Section 12.9 closed chapter 12 by naming the three mechanisms this volume's central impossibility theorems rest on — theorem 2.1's pointwise curvature mismatch, chapter 6's topological index argument, and theorems 11.8 and 12.1's global area/angle bound — and observing that all three are, in one way or another, obstructions built from second-order data. Theorem 13.2 is this volume's first result of the opposite character: not a fourth obstruction, but a demonstration that the first three, however they are individually derived, are all specifically obstructions to *rigid* (at least $C^2$) geometry, with essentially no purchase at all one derivative down. Lemma 13.9's proof that no compact $C^2$ surface is flat, set directly against section 13.7's report that the identical flat metric *is* realized, exactly and everywhere, by a $C^1$ embedding, makes the volume's own title more precise than its four central theorems alone could: what this volume has shown impossible, throughout, is not the coexistence of a given metric with $\mathbb{R}^3$ as such, but the coexistence of that metric with $\mathbb{R}^3$ *smoothly enough to speak of curvature at all.* Loosen that one requirement, and — for the flat torus, and by theorem 13.2's full generality for any short-embeddable Riemannian manifold whatsoever, regardless of curvature sign, genus, or completeness — the obstruction this volume has spent four chapters establishing simply does not apply.

**Todo — verify primary source**

This chapter's own open items, gathered in one place: only Kuiper's second 1955 paper (pages 683–689, on non-compact manifolds) remains at citation-only standing among this chapter's four named primary sources; Nash's 1954 paper and Kuiper's Part I are now read via Poznyak's 1957 translation (sections 13.5 and 13.6, with section 13.6.1's standing caveat), and the Borrelli–Jabrane–Lazarus–Thibert 2012 paper is now read in full in its published original (section 13.7). Gromov's $h$-principle framework (remark 13.4) is named for context and not developed; this chapter makes no claim to have verified that theorem 13.2 is correctly classified as an instance of it beyond what is common knowledge in the field. Proposition 13.5 and section 13.4 and lemma 13.9 and corollary 13.11 are this chapter's own independently checked contributions — the first computational, the second a from-scratch proof — and are not qualified by any of the above.

# Part III

# Historical Case Studies

# Chapter 14

# Mercator and Lambert: Two Solutions, and an Eighteenth-Century Sighting of the Theorem

## 14.1 Why Part III opens with this pair, and not with something else

Part I derived, from the intrinsic geometry of the sphere alone, that no projection can be simultaneously conformal and equal-area (theorem 4.1), and Part II sharpened that impossibility into a global distortion bound (theorem 5.2, theorem 5.10) and a topological obstruction to fixed-direction fields (chapters 6 and 7). Part III's task is different in kind from Parts I–II, and it is worth being explicit about the change of register before beginning, in the same spirit as chapter 1's three-way distinction between empirical, historical, and impossibility claims. The theorems already proved are impossibility claims; nothing in this chapter or the two that follow it adds to them. What Part III adds is the historical record of how working cartographers, without access to Gauss's 1827 proof or in some cases without any proof at all, arrived independently at solutions occupying each of the three families that corollary 4.5 shows exhaust the possibilities — and, in one case treated below, arrived at a qualitative statement of the impossibility theorem itself, a century and a half before chapters 2 and 4 made it rigorous.

This chapter treats the two earliest and most consequential solutions: Mercator's conformal cylindrical projection of 1569, and Lambert's equal-area azimuthal projection of 1772. Between them they populate two of the three families of corollary 4.5 with the constructions that, more than any others, fixed those families' names in the cartographic literature; chapter 16 and chapter 18 carry the historical record forward to the aphylactic/compromise family and to the twentieth- and twenty-first-century discretization schemes documented in the DGGS companion volume.

## 14.2 Mercator, 1569: solving conformality without a theory of conformality

Gerardus Mercator published his world map in August 1569, under the title *Nova et aucta orbis terrae descriptio ad usum navigantium emendate accommodata* — a large planisphere,

printed on eighteen separate sheets, whose title already announces its purpose: a description of the Earth "corrected for the use of navigators."[1] The construction he arrived at is exactly the one this volume's chapters 2 and 3 would later show is the *unique* way (up to the choice of central meridian and scale) to satisfy a cylindrical layout — meridians as equally spaced parallel vertical lines, parallels as horizontal lines — while keeping $a(P) = b(P)$ at every point: the conformal cylindrical projection, whose parallel spacing must grow with latitude precisely as the integral of the secant function, is what is needed to keep a straight line of constant compass bearing on the map (a rhumb line, or loxodrome) corresponding to a straight line on the sphere's own loxodromic curves.

**Key Result**

*In the language of definition 4.4, Mercator's projection is a member of the conformal family: $\omega(P) \equiv 0$ everywhere it is defined (the open cylinder minus the poles), which by corollary 4.5 forces it to fail equivalence everywhere except along the single standard parallel (the equator, in the normal aspect Mercator used), where $s = 1$ exactly. The areal inflation away from that parallel — famously, Greenland rendered comparable in printed size to Africa, though the true ratio is roughly fourteen to one — is not an incidental defect of Mercator's specific formula; it is the unavoidable cost, quantified exactly by proposition 4.3's trade-off curve at $\omega = 0$, of the property he was actually optimizing for, which was not areal fidelity but the straight-rhumb-line property that follows from conformality together with the specific cylindrical layout.*

What makes the 1569 map historically remarkable for this volume's purposes, rather than merely as a practical achievement, is what Mercator did *not* supply. He gave no derivation of the parallel spacing, no stated formula, and no proof that the construction he used was the only one, or even a distinguished one, among constructions with the straight-rhumb-line property.[2] The gap between Mercator's practical solution and a theoretical justification of it is, in miniature, the same gap this volume's whole method is designed to close for the case of impossibility rather than existence: knowing that a construction works, or does not, is a different and logically prior question from knowing *why*, and cartographic history repeatedly produced the former well before anyone supplied the latter.

**Todo — verify primary source**

The account above rests entirely on tertiary and secondary sources (general cartographic reference works and one academic reconstruction of the map's cartometric accuracy), not on Mercator's original map or legend text, which has not been consulted directly in this pass. The specific historical question of how Mercator computed his table of parallel spacings without a published derivation — whether by graphical interpolation from Portuguese portolan practice, from correspondence with Pedro Nunes, or by some other route — is reported above as *unsettled in the consulted secondary literature*, not resolved; a direct reading of Mercator scholarship (Nordenskiöld 1889, or the relevant chapter of Snyder's *Flattening the Earth*, already flagged as inaccessible in full to this series in the qibla volume's preface) would be needed before this volume could go further. Section 14.3 narrows, but does not close, this item.

[1] Title and 1569 date confirmed against multiple independent secondary accounts of the map's publication and dedication (to Duke Wilhelm of Cleves), consulted 1 Aug 2026: the Mercator projection and Mercator 1569 world map entries maintained collaboratively as tertiary reference sources, an academic reconstruction of the map's projection accuracy by Snyder-tradition cartometric analysis, and a historical account of the map's construction and reception. This chapter has not consulted the map itself, nor Mercator's original legend text, directly; the todobox at the end of this section records that limitation explicitly, in keeping with this volume's sourcing discipline.

[2] This absence is itself well documented in the secondary literature on the map's construction and reception; several later authors, including Edward Wright in 1599, worked out the mathematical justification — the integral of the secant, and its logarithmic-tangent closed form — that Mercator's own published map never states, and how Mercator himself arrived at the correct table of parallel spacings without publishing that derivation remains, per these accounts, not fully settled.

## 14.3 Legend 3 of the 1569 map: an English translation of Mercator's own stated method, reached at one remove

A later pass located an actual route into Mercator's own words, rather than into secondary paraphrase of them: the fifteen framed legends of the 1569 map were transcribed and translated into English by the International Hydrographic Bureau in 1932, published as "Text and Translation of the Legends of the Original Chart of the World by Gerhard Mercator, issued in 1569" in the *Hydrographic Review*, volume 9, number 2, pages 7–45; that 1932 translation, reproduced with the Bureau's permission on a collaboratively maintained tertiary reference page, was reached and read in this pass (consulted 1 Aug 2026).[3]

**Key Result**

*[Legend 3 states the method in Mercator's own words, and confirms he left the derivation unstated] Legend 3, addressed by Mercator to the reader under the heading* Inspectori Salutem*, gives three stated aims for the map, the first of which is the construction this chapter treats: an arrangement in which position, direction, and distance correspond correctly between sphere and plane, achieved, in Mercator's own stated words as translated in 1932, by having "progressively increased the degrees of latitude towards each pole in proportion to the lengthening of the parallels with reference to the equator." This is a correct qualitative description of exactly the secant-integral spacing section 14.2 already derives from first principles; what Legend 3 does not supply, any more than the secondary accounts already consulted reported, is a stated formula, a derivation, or a description of the method by which the table of spacings was actually computed. On this specific point — the absence of a derivation — the primary legend text and the secondary literature already cited now agree exactly, which upgrades that absence from a reported characterization to a directly confirmed one.*

**Todo — verify primary source**

Section 14.3 closes the narrow question of whether Mercator's own text states a formula or derivation (it does not, confirmed directly rather than via secondary report), but leaves open the question section 14.2's original todobox posed: by what actual procedure the table of parallel spacings was computed. Legend 3 itself is silent on this; nothing in the pages read here bears on the graphical-interpolation, portolan-practice, or Nunes-correspondence hypotheses reported, unresolved, in the secondary literature. What remains unread is the Latin original as an independent check on the 1932 translation, the Bureau's own *Hydrographic Review* archive rather than the tertiary mirror used here, and the fourteen legends beyond Legend 3, several of which (Legends 10 and 12, on the Organum Directorium and on rhumb-line distance) bear on this volume's subject matter but were not the target of this pass and so were not read with the same care.

## 14.4 Lambert, 1772: the equal-area solution, and a qualitative statement of theorem 4.1 itself

Johann Heinrich Lambert's 1772 memoir *Anmerkungen und Zusätze zur Entwerfung der Land- und Himmelscharten* ("Notes and additions to the design of terrestrial and celestial charts"), published as part of the third volume of his *Beyträge zum Gebrauche der Mathematik und deren Anwendung*, introduced seven new projections in a single publication, unnamed by Lambert himself but individually identified and named by later cartographers; among them are the projection now called the Lambert azimuthal equal-area projection and

[3] This is a translation of Mercator's original Latin legend text, not the Latin itself, and it is reached through a tertiary mirror of the 1932 *Hydrographic Review* translation rather than through the Review's own archive directly; both removes are recorded here explicitly, in keeping with this volume's practice of distinguishing a text from an account of a text. The Latin original of Legend 3, also reproduced on the same mirror page, was not separately checked word-for-word against this volume's own reading of the English translation.

the Lambert conformal conic.[4] The azimuthal equal-area construction is a member of the equal-area family in exactly the sense of definition 4.4: $s(P) \equiv 1$ everywhere the projection is defined (the whole sphere except the point antipodal to the projection's centre), which forces, by corollary 4.5, that it cannot be conformal anywhere except at the single centre point itself, where $a = b = 1$ trivially. The distortion this projection cannot avoid is angular rather than areal: shapes near the projection's rim are visibly sheared, a fact any modern reader can verify directly from a printed world map in this projection, and one whose exact quantitative form, at any given angular distance $\rho$ from the centre, is again read off directly from proposition 4.3's curve at $s = 1$.

What earns Lambert's 1772 memoir its place in this chapter as more than a second named construction, however, is a claim made about the memoir in the secondary cartographic literature and independently corroborated across sources: that Lambert, in the course of presenting his seven projections, explicitly noted that conformality and equivalence are properties a single projection cannot hold simultaneously.[5]

**Key Result**

*If the secondary-source attribution above is accurate, Lambert's 1772 memoir contains what this chapter proposes to call the volume's earliest recorded sighting of theorem 4.1: a correct qualitative statement of the theorem this volume proves rigorously in chapter 4, arrived at fifty-five years before the differential-geometric machinery (Gauss's Theorema Egregium) needed to prove it existed, and presumably arrived at by direct inspection of his own seven constructions — observing empirically that his conformal conic and his equal-area azimuthal could not be reconciled into one formula — rather than by a general argument covering every possible projection. This is exactly the distinction section 14.2 drew for Mercator's rhumb-line construction, run in the opposite direction: where Mercator solved a case without articulating the general principle behind it, Lambert appears to have articulated a true general principle two centuries before this volume's proof of it, without the geometric apparatus to prove it himself.*

**Todo — verify primary source**

The claim that Lambert explicitly stated the conformal/equal-area exclusivity rests on two secondary sources agreeing with each other, not on a direct reading of the 1772 *Anmerkungen* (in the original German or in Tobler's 1972 English translation, *Notes and Comments on the Composition of Terrestrial and Celestial Maps*, neither consulted in this pass), so the exact form of Lambert's statement — whether it is a general remark about projections as a class, or a narrower observation confined to the seven constructions he presents — is not yet established to this volume's usual standard and is reported above with that qualification attached explicitly. Locating and reading the relevant passage of the Tobler translation, which is the standard access point for this text in the English-language secondary literature, is the natural next step and is left open here rather than resolved provisionally.

[4] Publication details (third part of the *Beyträge*, Berlin 1772, the chapter titled as above) and the count of seven new projections are corroborated across several independent secondary sources consulted 1 Aug 2026, including a peer-reviewed historical account of Lambert's cartographic work and a separate cartometric paper on the equal-area azimuthal projection specifically; both cite the 1894 Wangerin critical reprint (Ostwald's Klassiker series) as the standard modern edition. The original 1772 German text and the 1972 Tobler English translation have not been consulted directly in this pass; see the todobox below.

[5] This specific historical claim is reported, in near-identical language, by two secondary sources with no evident dependence on one another: a 2025 comparative paper on Lambert's derivation of the Mercator projection as a special case of his conformal-conic construction, and a separate paper on the Lambert equivalent azimuthal projection specifically, which attributes the observation (via an intermediate historical account by Frischauf) to Lambert's own 1772 text. Both describe this as Lambert pointing to the mutual exclusivity of the two properties, not as Lambert proving it in the sense of chapters 2 and 4: no source consulted here claims Lambert supplied anything resembling a Gaussian-curvature argument, which would have required results (the Theorema Egregium itself) not published until 1827, fifty-five years later and after Lambert's own death in 1777.

## 14.5 The primary text read directly, and a narrowing rather than a confirmation

A later pass located and read the 1772 memoir itself, not in Tobler's translation but in the same 1894 Wangerin critical edition (*Ostwald's Klassiker der exakten Wissenschaften*, Nr. 54) both secondary sources cited above already treat as the standard modern edition. A Google-digitized scan hosted on the Internet Archive (identifier `anmerkungenundz00lambgoog`) proved fully machine-readable — unlike the scanned 1867/1869 material met elsewhere in this volume (sections 16.3 and 22.12), Wangerin's 1894 Antiqua-set German typesetting OCRs cleanly — and at 95 pages was retrieved and read in its entirety in a single pass (2 Aug 2026), from Lambert's opening five desiderata for a perfect map (§1) through the spheroidal-Earth extension of his general conformal construction (§123) that closes the memoir.

**Key Result**

*[What the primary text does and does not contain] The memoir's opening section states five conditions a perfect map would satisfy (undistorted shape, true area ratios, true distance ratios, great circles as straight lines, easy recovery of latitude and longitude) and immediately observes that a curved surface cannot meet them all on a flat sheet, so "man muss sich eine oder einige davon besonders vorsetzen" (one must single out one or a few of them to satisfy especially) — a general statement that a perfect map is impossible, not a statement that conformality and equal-area specifically are the pair in tension. §3, after surveying the three classical perspective projections (orthographic, stereographic, central), states that none of the three satisfies all the conditions of §1, and singles out the true-area-ratio condition as the one none of them meets. Sections IV (§48–57) and IX (§100–110) of the memoir — Lambert's general method for constructions preserving all angles, parametrized by an exponent $m$ that interpolates continuously between Mercator's cylindrical projection ($m = 0$) and the stereographic ($m = 1$) and yields the construction later named the Lambert conformal conic for other values of $m$, and his separate method for constructions preserving area — are presented as two independent projects addressing two different desiderata from §1's list, each solved in its own right; at no point in the sections read does Lambert state, as a general remark either about these two constructions or about projections as a class, that a single construction cannot do both. No search of this scan for Lambert's own name for either property, nor for a passage explicitly juxtaposing the conformal and equal-area constructions as mutually exclusive, located such a statement.*

This narrows section 14.4's open item without closing it outright, and in a direction opposite to what a confirming read would have given: the primary text, read across its full 95 pages, supports the weaker and more general claim already in §1 and §3 (no projection meets every desideratum; area-truth is the specific one Lambert flags as failed by the three perspective constructions he starts from) but does not, in the passages read here, supply the sharper conformal-versus-equal-area statement the two secondary sources attribute to Lambert via Frischauf's intermediate account. Three qualifications keep this from being treated as a refutation. First, Wangerin's edition reprints the 1772 text essentially verbatim (per the editor's own prefatory note) but this pass read it in translation-free German prose dense with eighteenth-century mathematical notation, so a short remark could in principle have been misread as routine derivation rather than recognized as the specific claim being sought; a second, targeted reading aimed explicitly at this question, rather than a single sequential pass, would be a stronger test. Second, Frischauf's own intermediate account (cited by the secondary sources but itself unread here) may draw on a different edition, a different treatise of Lambert's, or a inference from the juxtaposition of Lambert's results rather than a claim that Lambert stated the exclusivity in so many words — in which case both this pass and the secondary sources could be correct about what they each examined, and the mismatch would be an artifact of what counts as "stating" the exclusivity rather than

a factual disagreement. Third, the memoir's editorial apparatus (Wangerin's own appended notes, read only in part here) was not searched as thoroughly as Lambert's own text for a relevant remark. Section 14.5 is accordingly reported as a narrowing of the open item, not a closure of it: the primary text is now read in full for the first time in this volume, and what it directly supports is weaker than what the secondary-source attribution claimed.

Corollary 4.5 asserted that every projection belongs to exactly one of three mutually exclusive families along the conformal/equal-area axis, with the aphylactic family as the residual case. Mercator and Lambert, read together, are the two oldest and most historically consequential occupants of the two non-residual families, and table 14.1 records the placement explicitly, extending the seed classification section 4.4 began in chapter 4.

| Projection | Author, date | Family (definition 4.4) | Exact locus ($\omega = 0$ or $s = 1$) |
|---|---|---|---|
| Cylindrical conformal | Mercator, 1569 | Conformal | entire domain ($\omega \equiv 0$) |
| Azimuthal equal-area | Lambert, 1772 | Equal-area | entire domain ($s \equiv 1$) |

Table 14.1: The two constructions treated in this chapter, placed in the taxonomy of definition 4.4 and corollary 4.5. Both are global occupants of their family (the defining property holds on the whole domain, not merely along an exceptional curve), which is what makes them the canonical named representative of each family in the wider cartographic literature and in the companion volume *No Map Is Best*.

**Epistemic Notice**

The specific numerical distortion figures for Mercator's and Lambert's projections at named latitudes or angular distances — the areal inflation factor at 60°N on Mercator's map, the maximum angular shear near the rim of Lambert's azimuthal disk — are computed and tabulated in *No Map Is Best*. They are not reproduced in table 14.1 above, which records only the qualitative family placement provable from corollary 4.5 directly; the companion volume's specific figures are recalled from memory of that volume's contents, not re-read from its source files in this pass, and are therefore deliberately omitted here rather than restated at a precision this chapter cannot presently verify. A future revision that re-reads the relevant chapter of *No Map Is Best* directly could add a third column of exact figures to table 14.1 with full sourcing; this is left as an explicit open item rather than done provisionally from recollection.

Two of the three named case studies originally sketched for Part III — the transverse/secant refinement represented by the Gauss–Krüger transverse Mercator, and the twentieth-century discrete grid systems documented at length in the DGGS companion monograph — remain to be drafted, together with chapter 26's closing two-axis synthesis anticipated in section 4.4's disclaimer. Chapter 16 takes up the first of these: a case chosen specifically because, unlike the global constructions of this chapter, it exhibits the *regional* strategy of accepting global distortion in exchange for near-zero local distortion along a chosen line, connecting Part III's historical record forward to chapter 3's local machinery in a way sections 14.2 and 14.4 above, being both single-projection global constructions, did not have occasion to illustrate.

# Chapter 15

# Loxodromes: The Curve Mercator Was Actually Built to Straighten

## 15.1 A complement, not a gap

Chapter 14 treats the Mercator projection as a case study in the conformal/equal-area taxonomy of chapter 4: conformal, cylindrical, not equal-area. This is the correct level of generality for that chapter's purpose, but it leaves unstated the actual reason Mercator's 1569 chart was constructed at all — not angle-preservation in the abstract, but a specific navigational curve that angle-preservation makes into a straight line. This short chapter is explicitly a complement to chapter 14, not a new structural gap in the volume's impossibility framework; it is filed here, immediately after chapter 14, on exactly that footing.

**Epistemic Notice**

This chapter's content is standard 16th–20th century navigational geometry, available in any treatment of the Mercator projection; no primary 1569 source has been consulted, and none of the results here are new. The chapter's only purpose is to make explicit, with a full derivation from this volume's own machinery, a fact chapter 14 uses implicitly.

## 15.2 Definition and the differential equation

**Definition 15.1.** A *loxodrome* (or *rhumb line*) on $S^2_R$ is a curve that crosses every meridian at a fixed angle $\beta \in (0, \pi/2]$ to the local north direction, $\beta$ constant along the curve. (The degenerate case $\beta = 0$ is a meridian itself; $\beta = \pi/2$ is the equator or, more generally, a parallel is *not* a loxodrome unless it is the equator, since a parallel does not cross meridians transversally in the relevant sense — see remark 15.5 below.)

Parametrize the sphere by latitude $\varphi$ and longitude $\lambda$ as in chapters 2 and 3. Along any curve $\lambda(\varphi)$, the local infinitesimal displacement has a north-south component $R\,d\varphi$ and an east-west component $R\cos\varphi\,d\lambda$ (the standard spherical arc-length elements, already used in chapter 14's construction). The condition of crossing every meridian at constant angle $\beta$ to north is

$$\tan\beta = \frac{R\cos\varphi\,d\lambda}{R\,d\varphi} = \cos\varphi\,\frac{d\lambda}{d\varphi}.$$

**Proposition 15.2** (Loxodrome equation)**.** *A curve $\lambda(\varphi)$ is a loxodrome of constant bearing $\beta$ if and only if*

$$\lambda(\varphi) - \lambda_0 = \tan\beta \int_{\varphi_0}^{\varphi} \sec\varphi'\,d\varphi' = \tan\beta\,\Big[\psi(\varphi) - \psi(\varphi_0)\Big],$$

*where* $\psi(\varphi) = \ln\tan(\pi/4 + \varphi/2)$ *is the spherical isometric latitude already introduced in chapter 14 (and generalized to the ellipsoid in eq.* (17.1) *of chapter 17).*

*Proof.* Separate variables in the defining equation: $d\lambda = \tan\beta \sec\varphi\, d\varphi$, and integrate, using $\int \sec\varphi\, d\varphi = \ln\tan(\pi/4 + \varphi/2) = \psi(\varphi)$ (a standard elementary integral, checked by differentiating the right-hand side). □

**Key Result**

*Proposition 15.2 says a loxodrome is exactly a straight line,* $\lambda - \lambda_0 = \tan\beta\,(\psi - \psi_0)$*, in the coordinates* $(\psi, \lambda)$ *— which are, by definition, the Mercator plotting coordinates of chapter 14 up to the constant scale factor* $R$*. This is the entire content of Mercator's construction: a projection was sought, and found, in which this one family of navigationally meaningful curves becomes straight, at the cost (by chapter 4) of every area on the map being wrong away from the equator.*

## 15.3 The loxodromic spiral and its infinite length

**Proposition 15.3** (Loxodromes spiral into the poles)**.** *For* $\beta \in (0, \pi/2)$*, a loxodrome starting at any* $\varphi_0 < \pi/2$ *approaches the pole* $\varphi = \pi/2$ *only after crossing every meridian infinitely often, and its arc length from* $\varphi_0$ *to the pole is finite.*

*Proof.* As $\varphi \to \pi/2^-$, $\psi(\varphi) = \ln\tan(\pi/4 + \varphi/2) \to +\infty$, so by proposition 15.2, $\lambda(\varphi) - \lambda_0 = \tan\beta \cdot [\psi(\varphi) - \psi_0] \to \infty$: the longitude winds around infinitely many times, which is the spiraling claim. For arc length, the element is $ds = R\,d\varphi/\cos\beta$ (immediate from $\tan\beta = R\cos\varphi\, d\lambda/(R\,d\varphi)$ together with $ds^2 = R^2 d\varphi^2 + R^2\cos^2\varphi\, d\lambda^2 = R^2 d\varphi^2 \sec^2\beta$), so total length is $\int_{\varphi_0}^{\pi/2} R\sec\beta\, d\varphi = R\sec\beta\,(\pi/2 - \varphi_0)$, finite. □

*Remark* 15.4. This is a small, self-contained illustration of a phenomenon that recurs at much greater technical depth in chapter 25: a curve (there, a discrete curvature distribution) can wind around a singular point — here, literally — infinitely often while still being globally finite in the relevant measure (there, total angular defect; here, arc length). The loxodromic spiral is the cartographic series' simplest instance of that pattern.

*Remark* 15.5 (Loxodromes are not geodesics). Except along a meridian ($\beta = 0$) or the equator ($\beta = \pi/2$, where the loxodrome coincides with a great circle), a loxodrome of constant bearing $\beta \in (0, \pi/2)$ is *not* a great circle: a great circle's bearing relative to north changes continuously along its length (this is exactly the retroazimuthal-versus-conformal tension examined in full generality in chapter 7), while a loxodrome's bearing is constant by definition. Confusing the two — steering a constant compass bearing versus following the shortest path — is precisely the practical distinction that made great-circle sailing charts and loxodromic (Mercator) charts different tools for different purposes, and is the navigational content underneath chapter 7's abstract impossibility statement.

# Chapter 16

# Gauss, Krüger, and the Regional Strategy: Trading Global Coverage for Local Fidelity

## 16.1 A different strategy, not a third family

Chapter 14 treated two constructions that are each, in the sense of definition 4.4, a *global* occupant of their family: Mercator's projection is conformal on its entire domain, Lambert's azimuthal is equal-area on its entire domain, and table 14.1 recorded exactly that. This chapter's case is different in a way that is easy to miscategorize: the Gauss–Krüger projection (better known outside its native geodetic literature as the ellipsoidal transverse Mercator, and as the basis of the Universal Transverse Mercator, UTM, grid) is, like Mercator's own projection, a member of the conformal family in the strict sense of definition 4.4 — $\omega(P) \equiv 0$ everywhere it is defined. What changes is not the family but the *aspect and the intended use*: where Mercator's cylinder is tangent along the equator and the construction is meant to cover the whole navigable ocean at once, Gauss–Krüger's cylinder is tangent along a chosen *meridian*, and the construction is meant to be used only in a narrow longitudinal band around that meridian, discarded and re-centred on a new meridian once the user moves far enough east or west. This is, this chapter argues, the first fully worked historical instance of a strategy chapter 4 identified only abstractly: proposition 4.6's "curve of true scale," realized here as an explicit, engineered design choice rather than an incidental feature of a global formula.

## 16.2 From Lambert's spherical sketch to Gauss's ellipsoidal survey mathematics

The transverse aspect of the conformal cylindrical projection was not new when Gauss took it up. Lambert's 1772 memoir, already the subject of section 14.4 above for its azimuthal equal-area construction, is reported by the secondary geodetic literature to contain, among its seven projections, a spherical transverse Mercator as well — the same publication, on this account, seeding both the equal-area case study of section 14.4 and the conformal regional strategy taken up in this chapter, which this volume had no occasion to flag when section 14.4 was written and records here as a genuine point of contact between the two sections.[1] What

[1]This attribution is reported consistently across several independent secondary geodetic sources consulted 1 Aug 2026, including a conference paper on the Gauss–Krüger projection's history and a separate technical note on the Karney–Krüger equations, both citing Lambert's spherical transverse construction as the immediate

Lambert's spherical construction lacked, and what made it unsuitable for the high-precision national land surveys that were beginning across Europe in the early nineteenth century, was any accounting for the Earth's ellipsoidal (rather than spherical) shape — a correction whose absence matters enormously at the sub-metre precision a cadastral or geodetic survey requires, even though it is invisible at the scale of Mercator's own navigational chart.

Carl Friedrich Gauss supplied that correction in 1822, in the course of his work as director of the geodetic survey of the Kingdom of Hanover — the same practical occupation chapter 1 already noted as the historical seed of the Theorema Egregium itself, so that the same decade of Gauss's career that produced the volume's foundational Part I theorem also produced, on the secondary record, the specific projection this chapter treats.[2] Gauss's own published record of the 1822 construction is reported to be thin: several independent secondary accounts converge on the specific claim that Gauss left few details of his method, communicating the essential result (a conformal "double projection," first mapping the ellipsoid conformally onto a sphere, then applying the ordinary spherical transverse Mercator to that sphere) without a full derivation of the series expansions needed to compute it in practice.

**Key Result**

*The gap between a construction's existence and its derivation, already noted for Mercator in section 14.2, recurs here in an even more consequential form: it took two further, independently datable stages of secondary work — Oskar Schreiber's analysis of Gauss's method, published in 1866 and again in 1897, and Louis Krüger's 1912 monograph* Konforme Abbildung des Erdellipsoids in der Ebene *("Conformal mapping of the terrestrial ellipsoid onto the plane"), which re-evaluated both Gauss's and Schreiber's work and produced the direct ellipsoid-to-plane series expansion in practical use ever since — before Gauss's 1822 result became a fully specified, computable procedure. This is why the projection carries both names: "Gauss" for the 1822 conformal double-projection idea, "Krüger" for the 1912 computational form that made it usable, with an intermediate "Gauss–Schreiber" variant (used for the Prussian land survey of 1876–1923, with non-constant scale along the central meridian) attesting to the ninety-year gap between the two.*

**Todo — verify primary source**

This section's account rests entirely on secondary geodetic literature (a 2010s technical conference paper on the Gauss–Krüger projection's derivation and history, an ICA map-projections reference wiki entry, and a further technical note on the modern Karney–Krüger nanometre-accuracy equations), not on Gauss's 1822 correspondence or manuscripts, Schreiber's 1866/1897 papers, or Krüger's 1912 *Konforme Abbildung des Erdellipsoids in der Ebene* directly, none of which has been consulted in this pass. The claim that "Gauss left few details" is itself a secondary characterization reported consistently across sources rather than a claim this chapter can independently confirm by inspecting what Gauss did or did not publish.

predecessor of Gauss's ellipsoidal one; neither source, nor this chapter, has verified the claim against Lambert's own 1772 text, which section 14.4's todobox already flagged as unconsulted for the azimuthal case and which this chapter inherits as an open item for the transverse case as well.

[2]The dating (1822) and the connection to the Hanover survey are corroborated across the geodetic secondary literature cited throughout this section. As with chapter 2's account of the 1827 Disquisitiones, this reflects the well-documented fact that Gauss's differential-geometric and cartographic work of the 1820s arose from a single sustained practical engagement with the survey, not two unrelated projects; it does not imply this chapter has independently verified a direct textual link between the 1822 projection work and the 1827 memoir beyond what the secondary sources report.

## 16.3 Krüger's own monograph: an open-access digitization located, but not yet readable

A later pass, prompted specifically by the previous todobox's flagged gap, searched for Krüger's 1912 monograph directly rather than for secondary accounts of it, and found more than the usual antiquarian catalogue record: the GFZ German Research Centre for Geosciences (successor institution to the Prussian Geodetic Institute Krüger himself directed from 1917) hosts its own digitization of the work, with a persistent DOI (`10.2312/GFZ.b103-krueger28`) and an explicit unrestricted-access rights statement, cross-listed in the Deutsche Digitale Bibliothek and confirmed independently by GFZ's own publication repository (all three consulted 1 Aug 2026). This is a materially better access route than appendix A.7's general remark, made for the volume's other unconsulted primary sources, that accessibility alone had been checked without success in obtaining a text.

**Key Result**

*[The digitization is reachable, and its own running header cross-checks the antiquarian page count] The PDF at GFZ's own digitization address (`bib.gfz-potsdam.de/pub/digi/krueger2.pdf`) was fetched directly in this pass and returned content — not an access-denied response — confirming the file is genuinely open, consistent with the Deutsche Digitale Bibliothek's own rights statement ("Der Zugriff auf das Objekt ist unbeschränkt möglich"). Every page carries the identical running header "L. Krüger: Konforme Abbildung des…" and the same DOI, repeated once per page; the number of times this header recurs in the fetched text is consistent with the independently sourced antiquarian bookseller's collation note for the original 1912 Teubner printing ("IX, 172 S.," i.e. nine roman-numbered front-matter pages plus 172 numbered body pages), which this pass had not previously had reason to connect to a page count obtainable from the digitization itself.*

**Todo — verify primary source**

Section 16.3 narrows, but does not close, this chapter's Krüger item: the file is confirmed genuinely open-access and its extent cross-checked against an independent secondary collation, but the automated text extraction available in this pass returned only the repeating running header on each page, not the body text underneath it — consistent with a 1912 scan of a mathematically dense, equation-heavy typeset page (Fraktur-adjacent German mathematical typesetting of the period is a poor match for automated OCR, and GFZ's own digitization appears not to have added a body-text layer beyond the page headers) rather than with the file being inaccessible. This is a genuinely different failure mode from appendix A.5's HathiTrust and JSTOR blocks or appendix A.7's unconfirmed Goode lead: the document is open and was retrieved, but this pass's tooling could not read what is almost certainly a scanned image of Krüger's actual derivation underneath. A future pass with page-image inspection (rather than text extraction alone) would be the natural next step, and would close this item outright if it succeeds; this pass did not attempt it.

## 16.4 The central meridian as proposition 4.6 made concrete

Chapter 4 proved, conditionally on a regular-value hypothesis, that a projection failing conformality or equivalence on all of its domain can still achieve one property exactly along a distinguished curve (proposition 4.6), and flagged explicitly that it had not verified this regularity condition for any specific named projection in the companion series (section 4.4's closing disclaimerbox). The Gauss–Krüger construction supplies exactly such a verification, and does so by design rather than by incidental geometric accident.

**Key Result**

*In the basic (tangent) form of Gauss–Krüger, the central meridian is a curve of exact scale:* $\omega(P) = 0$ *holds identically across the whole domain by conformality, and additionally the linear scale factor* $k(P)$ *— not itself tracked by the* $(\omega, s)$ *pair of proposition 4.3, which records only the shape and area distortion at a point, but a further quantity derivable from the same principal-scale data of theorem 3.1 — equals exactly* 1 *along the central meridian and grows monotonically with distance from it. In the* secant *form used by UTM (scale factor* 0.9996 *at the central meridian itself), the true-scale locus splits into* two *parallel curves of exact scale, symmetric about the central meridian, exactly realizing proposition 4.6's abstract "exceptional set is a curve" conclusion twice over within a single named, operationally deployed projection. Where section 4.4 could point only to conic standard parallels as a plausible instance and left the regularity hypothesis unverified for any of them, the transverse cylindrical case is* $\nabla\omega \equiv 0$ *trivially satisfied (conformality holds everywhere, not just on the curve) and the scale-factor gradient across the central meridian is bounded away from zero except at the poles of the projection's own coordinate singularity, so the regularity condition needed for proposition 4.6's curve conclusion is not merely plausible here but follows from the explicit closed-form series for* $k(P)$ *that the Gauss–Krüger/Karney literature provides.*

This is the sense in which Gauss–Krüger is a genuinely different *strategy* from Mercator's, not merely a different aspect of the same one, and why this chapter's title separates "regional" from "global" rather than reusing chapter 14's frame with the cylinder rotated ninety degrees. Mercator's construction is used as published, over its whole domain, with distortion accepted as a smoothly increasing function of distance from a single fixed standard parallel. Gauss–Krüger, and every zoned coordinate system built on it, is used *only* in a narrow band — a UTM zone is typically 6° of longitude wide, so that even the zone's worst-case scale distortion at its eastern or western edge remains small — and then discarded in favour of a fresh copy of the same formula recentred on a new meridian for the next zone east or west. The projection formula never changes; what changes, zone to zone, is only which meridian plays the role of the exact-scale curve identified above. This is, in the vocabulary chapter 1 used to organize this volume's non-impossibility material, an engineering response to theorem 4.1 and to proposition 4.3's trade-off curve: since no single global formula can avoid growing distortion with distance from wherever exactness is placed, the practical solution adopted throughout twentieth-century national and international geodesy is to move where exactness is placed, repeatedly, rather than to seek a global formula exempt from the trade-off, which chapters 4 and 5 already show cannot exist.

**Epistemic Notice**

This chapter deliberately stops short of citing Monograph 5's DGGS material or *The New Datum*'s NATRF2022/NAPGD2022 account of the specific zone systems (State Plane, UTM, national Gauss–Krüger grids) currently used by real geodetic authorities, beyond the general remark above that they exist and are built on this projection. Those companion volumes contain the operational detail (exact zone boundaries, scale-factor conventions, adoption dates by country) that would substantiate the remark with citable specifics; recalling that detail from memory rather than re-reading the companion volumes' actual chapters would violate this series' sourcing discipline, so it is left as a cross-reference rather than restated here.

## 16.5 Placement and what Part III still owes

Table 16.1 extends table 14.1 with this chapter's case, and separates, for the first time in this volume, the axis section 4.4 used (conformal/equal-area/aphylactic family) from a second axis — global versus regional deployment — that chapter 4's closing disclaimerbox already anticipated as necessary for the eventual two-axis taxonomy of chapter 26, without

yet supplying an instance of the second axis in play. This chapter supplies that instance.

| Projection | Author, date | Family | Deployment |
|---|---|---|---|
| Cylindrical conformal | Mercator, 1569 | Conformal | Global, single instance |
| Azimuthal equal-area | Lambert, 1772 | Equal-area | Global, single instance |
| Transverse cylindrical conformal | Gauss 1822 / Krüger 1912 | Conformal | Regional, re-centred per zone |

Table 16.1: Extending table 14.1 with the regional-deployment axis this chapter introduces. The family column follows definition 4.4 exactly as before; the deployment column is new and anticipates the second axis of chapter 26's planned two-axis synthesis.

One case study remains before Part III's historical record is complete: chapter 18 turns from continuous, meridian-centred regional strategies to the twentieth- and twenty-first-century discrete global grid systems documented at length in the DGGS companion monograph, which reject the continuous-projection framework of every case treated so far in favour of a fundamentally different response to theorem 4.1 — not choosing where to place exactness within a single continuous formula, as Gauss–Krüger does, but discretizing the sphere directly and asking a differently shaped question about distortion, cell by cell rather than point by point. Only after that third case is in hand will Part IV's closing taxonomy (chapter 26) have a full set of historically grounded instances — one per family, one per deployment strategy, and one representing a genuinely different response to the impossibility theorems altogether — against which to check its claim of exhaustiveness.

# Chapter 17

# The Ellipsoidal Correction: Gauss–Krüger as It Is Actually Used

## 17.1 What the sphere-based chapter left out

Chapter 16 treats the Gauss–Krüger (transverse Mercator) construction on the sphere, which is the right level of generality for locating the projection within this volume's sphere-wide impossibility framework. But no national mapping agency, and no version of UTM (Universal Transverse Mercator) in practical use, works on a sphere. Gauss's own 1822–1825 Hanover survey work, cited already in chapter 2's historical note as the practical seed of the *Theorema Egregium* itself, was ellipsoidal from the start, and Krüger's 1912 completion of the closed-form conformal mapping — the very primary source flagged as still unobtained in chapter 16 and in the volume's preface — is a specifically ellipsoidal computation. This chapter supplies the piece chapter 16 deliberately did not attempt: what changes, mathematically, when the reference surface is the biaxial ellipsoid of revolution used for the real Earth (as in WGS84) rather than a sphere.

**Epistemic Notice**

This chapter's curvature and isometric-latitude formulas are standard results of ellipsoidal geodesy, reported here from Snyder, *Map Projections: A Working Manual* (1987), §3, not from a primary-verified reading of Krüger (1912) — the same unresolved primary source chapter 16 and the volume preface already flag. The chapter's logical content (that the impossibility framework of chapter 2 applies unchanged, and exactly what new machinery isometric latitude supplies) is original derivation from material already established in this volume.

## 17.2 The ellipsoid of revolution

Let $\Sigma_e \subset \mathbb{R}^3$ be the surface obtained by rotating the ellipse $\frac{x^2}{a^2} + \frac{z^2}{b^2} = 1$ $(a > b > 0)$ about the $z$-axis, parametrized by geodetic latitude $\varphi \in (-\pi/2, \pi/2)$ and longitude $\lambda \in [0, 2\pi)$ as

$$\mathbf{x}(\varphi, \lambda) = \big(N(\varphi)\cos\varphi\cos\lambda,\ N(\varphi)\cos\varphi\sin\lambda,\ N(\varphi)(1-e^2)\sin\varphi\big),$$

where $e^2 = 1 - b^2/a^2$ is the eccentricity and

$$N(\varphi) = \frac{a}{\sqrt{1 - e^2\sin^2\varphi}}$$

is the *prime vertical radius of curvature*, the distance from a point on $\Sigma_e$ to the axis of rotation measured along the surface normal. A second radius of curvature, the *meridian radius of curvature*

$$M(\varphi) = \frac{a(1-e^2)}{(1-e^2\sin^2\varphi)^{3/2}},$$

governs arc length along a meridian. Both are computable directly from the first and second fundamental forms of $\mathbf{x}(\varphi,\lambda)$ by the same route as chapter 2's §2–3; we omit the routine but lengthy computation of $E, F, G, e, f, g$ for $\Sigma_e$ here and quote the results, which agree with Snyder (1987), eqs. 4-18 to 4-20.

**Proposition 17.1** (Ellipsoidal Gaussian curvature)**.** *At geodetic latitude $\varphi$, the Gaussian curvature of $\Sigma_e$ is*

$$K(\varphi) = \frac{1}{M(\varphi)N(\varphi)} = \frac{(1-e^2)}{a^2(1-e^2\sin^2\varphi)^2}.$$

Two facts follow immediately from proposition 17.1 and are worth stating before anything else, because they settle the question this chapter's title implicitly raises — whether the ellipsoid needs its own impossibility theorem, separate from chapter 2's sphere-based argument.

**Corollary 17.2** (The ellipsoid inherits the sphere's impossibility unchanged)**.** *for every fixed eccentricity $e \in [0,1)$, $K(\varphi)$ is not only strictly positive at every $\varphi$ but bounded uniformly away from 0: since $0 \le e^2\sin^2\varphi \le e^2 < 1$, the denominator $(1-e^2\sin^2\varphi)^2$ is bounded above by 1, so $K(\varphi) \ge (1-e^2)/a^2 > 0$ for every $\varphi$. Consequently theorems 2.1 and 8.2 apply to $\Sigma_e$ exactly as stated for the sphere: no open region of $\Sigma_e$ admits an isometric flattening, for the same reason (strictly positive Gaussian curvature throughout), and chapter 4's conformal/equal-area exclusion holds on $\Sigma_e$ by the identical argument, since that proof (see chapter 4) uses only that $K$ is bounded away from 0 on the region in question, not that $K$ is constant.*

So the ellipsoid changes nothing about *whether* a conformal, equal-area, or fully distance-preserving map exists — it does not, for the same topological-curvature reason as the sphere. What it changes is the explicit *formula* for the best conformal map, because $K(\varphi)$ is no longer constant, and this is where Krüger's actual 1912 contribution lies.

## 17.3 Isometric latitude and the conformal mapping

The construction of a conformal (angle-preserving) map from $\Sigma_e$ to the plane proceeds, as it does for the sphere in chapter 14, by first finding coordinates in which the ellipsoid's metric is a scalar multiple of the flat metric ($E = G$, $F = 0$, the *isothermal* or *conformal* coordinate condition), and then applying an arbitrary holomorphic map of the resulting coordinate, which is automatically conformal by the Cauchy–Riemann argument reused throughout this volume (see chapter 22's from-scratch proof that composition of conformal maps is conformal). For the sphere this coordinate is the ordinary Mercator $y$-coordinate, $y = R\ln\tan(\pi/4+\varphi/2)$. For the ellipsoid, the corresponding *isometric latitude* is

$$\psi(\varphi) = \ln\left[\tan\left(\frac{\pi}{4}+\frac{\varphi}{2}\right)\right] - e\,\operatorname{artanh}(e\sin\varphi) = \ln\left[\tan\left(\frac{\pi}{4}+\frac{\varphi}{2}\right)\left(\frac{1-e\sin\varphi}{1+e\sin\varphi}\right)^{e/2}\right]. \tag{17.1}$$

**Proposition 17.3.** *In the coordinates $(\psi,\lambda)$ of eq. (17.1), the first fundamental form of $\Sigma_e$ is $ds^2 = N(\varphi)^2\cos^2\varphi\,(d\psi^2 + d\lambda^2)$: conformal to the flat plane, with $F = 0$ and equal diagonal coefficients.*

*Derivation of eq.* (17.1) *from* $E, F, G$. In the $(\varphi, \lambda)$ chart, $\Sigma_e$'s first fundamental form is $E = M(\varphi)^2$, $F = 0$, $G = N(\varphi)^2 \cos^2\varphi$ (already quoted above from Snyder). A coordinate change $\varphi \mapsto \psi(\varphi)$, $\lambda$ fixed, is the isothermal condition sought in proposition 17.3 exactly when it makes the diagonal coefficients equal after the substitution: writing $ds^2 = M^2 d\varphi^2 + N^2 \cos^2\varphi \, d\lambda^2$ and requiring $d\varphi = \dfrac{N\cos\varphi}{M} d\psi$ turns this into $ds^2 = N^2 \cos^2\varphi \, (d\psi^2 + d\lambda^2)$, proposition 17.3's claimed form, provided

$$\frac{d\psi}{d\varphi} = \frac{M(\varphi)}{N(\varphi)\cos\varphi}.$$

Substituting $M(\varphi) = a(1-e^2)/(1-e^2\sin^2\varphi)^{3/2}$ and $N(\varphi) = a/(1-e^2\sin^2\varphi)^{1/2}$ from proposition 17.1's derivation, the two powers of $a$ and two of $(1-e^2\sin^2\varphi)^{1/2}$ cancel, leaving the genuinely elementary integrand

$$\frac{d\psi}{d\varphi} = \frac{1-e^2}{(1-e^2\sin^2\varphi)\cos\varphi}.$$

That this integrates in closed form to eq. (17.1) — despite the meridian arc length $\int M(\varphi)\, d\varphi$ itself having no elementary closed form, the "genuine surprise" the previous flag noted — is confirmed directly rather than merely asserted: differentiating the right-hand side of eq. (17.1) with respect to $\varphi$ and simplifying reproduces the boxed integrand above term for term (partial-fraction decomposition of $1/[(1-e^2\sin^2\varphi)\cos\varphi]$ into a sec-type piece, integrating to the $\ln\tan(\pi/4+\varphi/2)$ term, and an $e\sin\varphi/(1-e^2\sin^2\varphi)$ piece, integrating to the $-e\,\mathrm{artanh}(e\sin\varphi)$ term, is the underlying mechanism), and the $e \to 0$ limit of the derivative reduces exactly to $\sec\varphi$, matching the spherical Mercator coordinate's own derivative and confirming the formula degenerates correctly to chapter 14's sphere case when eccentricity vanishes. □

Equation (17.1) is now derived from $E, F, G$ above rather than only quoted, closing the item the previous todobox flagged. The derivation's two nontrivial steps — that the isothermal-coordinate ODE $d\psi/d\varphi = M/(N\cos\varphi)$ reduces to an elementary integrand once $M, N$'s common $(1-e^2\sin^2\varphi)^{1/2}$ factor cancels, and that this integrand's antiderivative is exactly eq. (17.1) — were checked symbolically (differentiating the closed form and simplifying against the integrand) and numerically (evaluated at twelve random $(\varphi, e)$ pairs in $\varphi \in (-1.4, 1.4)$, $e \in [0, 0.3)$, all agreeing to within $2 \times 10^{-15}$, and at $\varphi = 0.5$, $e = 0.0818$, the WGS84-like value, agreeing to within $6 \times 10^{-17}$), and the $e \to 0$ spherical limit was checked to reduce exactly to $1/\cos\varphi$ rather than merely approximately. This is a from-scratch re-derivation of a formula previously only quoted from Snyder (1987); it does not independently verify Krüger's own 1912 route to the same result, which remains a separately flagged, still-unconsulted primary source (section 16.3).

Once proposition 17.3 is granted, the ellipsoidal Gauss–Krüger (transverse Mercator) construction proceeds by exactly the same complex-analytic route as the spherical case treated in chapter 16 — a conformal rotation of the sphere-analogue coordinate system to bring a chosen central meridian to the equator role, followed by the same Mercator-type unrolling — with $\psi$ standing in for the spherical isometric latitude throughout. The scale factor along the central meridian, denoted $k_0$ in every national UTM specification (e.g. $k_0 = 0.9996$ for standard UTM zones), and the meridian convergence angle $\gamma$ giving the difference between grid north and true north away from the central meridian, are both direct outputs of this same construction; we do not reproduce the full power-series expansion in $(\lambda - \lambda_0)$ used for practical computation, which runs to several pages in Snyder (1987) and adds no further impossibility-theoretic content beyond corollary 17.2.

## 17.4 Meridian arc length: the one place ellipsoidal geodesy has no closed form at all

One further disanalogy with the spherical case is worth recording, since it is the clearest illustration of what the ellipsoid costs computationally even though it costs nothing conceptually (corollary 17.2). On the sphere, arc length along a meridian from the equator to latitude $\varphi$ is simply $R\varphi$. On the ellipsoid, it is

$$m(\varphi) = \int_0^{\varphi} M(\varphi')\, d\varphi',$$

and $M(\varphi')$, given in closed form above, has no elementary antiderivative: $m(\varphi)$ is a complete elliptic integral of the second kind (up to reparametrization), not a finite combination of elementary functions of $\varphi$. Every practical implementation instead expands $m(\varphi)$ as a convergent trigonometric series in even multiples of $\varphi$ with coefficients in powers of $e^2$ (the standard six-term series is accurate to sub-millimeter precision for the Earth's actual eccentricity, $e^2 \approx 0.00669$). This is a genuinely different mathematical situation from anything in chapters 14 and 16: not an impossibility in the sense the rest of this volume proves impossibilities, but a plain transcendence obstruction, and worth flagging as such precisely because this volume is otherwise careful to distinguish "impossible" from "merely lacks a closed form."

*Remark* 17.4. This also explains, incidentally, why chapter 16's sphere-based treatment was a reasonable simplification for the purposes of *this* volume rather than an evasion: every impossibility-theoretic fact of interest (corollary 17.2) transfers unchanged from sphere to ellipsoid, while every fact that does *not* transfer unchanged (the isometric-latitude formula, the transcendence of meridian arc length) is a matter of explicit formula and numerical practice, not of what is or is not achievable.

# Chapter 18

# Discrete Global Grid Systems: Abandoning the Continuous Chart Altogether

## 18.1 A third response, not a third family

Chapters 14 and 16 treated three constructions — Mercator's, Lambert's, Gauss–Krüger's — that respond to theorem 4.1 in two ways this volume has now named explicitly: accept the trade-off globally along a single formula (chapter 14), or accept it locally by repeatedly relocating where the trade-off vanishes (chapter 16). Both strategies share an unstated premise: that the goal is a single continuous chart, $f : U \to \mathbb{R}^2$, covering as much of the sphere as the strategy allows. The case this chapter takes up abandons that premise altogether. A discrete global grid system (DGGS) does not seek a continuous chart of the sphere at all; it partitions the sphere directly into a hierarchy of finitely many cells and represents geographic data cell-by-cell, so that the object needing to be assessed for distortion is not a scale-and-angle field $(\omega(P), s(P))$ varying continuously over a domain, but a finite combinatorial structure — a polyhedral tiling — whose distortion is a property of each cell, not of a point.

This is the subject already surveyed at length, computationally and constructively, in the companion Monograph 5 (*From Continuous to Discrete*), and this chapter does not attempt to repeat that survey's scope. Its purpose here is narrower and specific to this volume's project: to place the DGGS strategy inside the same three-family, two-axis taxonomy chapters 14 and 16 have been building, and, in section 18.3 below, to draw out an unplanned connection between the combinatorial structure every DGGS is forced to contain and the topological obstruction already proved in this volume as chapter 6.

## 18.2 The founding survey: five design choices, one representative construction

The DGGS field's own account of its scope and internal structure is given by Sahr, White & Kimerling's 2003 survey, "Geodesic Discrete Global Grid Systems," published in *Cartography and Geographic Information Science* 30(2), pp. 121–134.[1] The survey's central organizing

[1]Citation details confirmed directly against the journal's own publication record (Taylor & Francis, DOI 10.1559/152304003100011090) and cross-checked against several independent secondary citations of the paper, consulted 1 Aug 2026. The paper's abstract and the five-way design-choice decomposition quoted in section 18.2 are drawn from the paper's own published abstract, not from a full reading of the paper's body

claim is that every system in this class — and the paper surveys a substantial number of them, spanning several decades of independent proposals — can be decomposed into exactly five substantially independent design choices: a base regular polyhedron (almost always the icosahedron among the systems surveyed, for reasons section 18.3 below returns to); a fixed orientation of that polyhedron relative to the Earth; a hierarchical partitioning method applied symmetrically within each face; a method for transferring that planar, face-level partition onto the curved sphere or ellipsoid; and a method for assigning a representative point to each resulting cell. Working through this decomposition, the survey constructs its own representative example, the Icosahedral Snyder Equal-Area aperture-3 hexagon (ISEA3H) DGGS, using an equal-area transfer method previously proposed by Snyder in 1992 for polyhedral globes generally.[2]

**Key Result**

*The choice of an* equal-area *transfer method in the Sahr–White–Kimerling representative construction is not incidental to this volume's concerns: it places the ISEA3H DGGS, cell by cell, in the equal-area family of definition 4.4 in exactly the sense section 14.4 placed Lambert's azimuthal projection there, except that the property now holds not as $s(P) \equiv 1$ for a continuously varying point $P$, but as an exact area match between each discrete cell and its corresponding fraction of the sphere's surface. Theorem 4.1 forbids a continuous projection from being both conformal and equal-area on any open set; nothing in that theorem, or in corollary 4.5's trichotomy, was stated for a discrete cell decomposition, and this chapter's task in section 18.3 is precisely to ask what, if anything, survives the passage from continuous domain to discrete cells.*

**Todo — verify primary source**

This section's account of the five-way design-choice decomposition and the ISEA3H construction is drawn from the Sahr–White–Kimerling (2003) paper's published abstract and its citation record in the DGGS literature, not from a full reading of the paper's body in this pass, and in particular the precise construction of the Icosahedral Snyder Equal Area transfer itself (the formula Snyder 1992 supplies) has not been reproduced or verified here. This is a narrower and more provisional treatment than Monograph 5 gives this same source material; readers wanting the fully verified account, including independent geometric reproduction of several DGGS constructions built directly from source code, should consult that companion volume rather than treat this chapter's brief placement as a substitute for it.

## 18.3 Twelve unavoidable defects: a discrete echo of chapter 6

Almost every DGGS surveyed by Sahr, White & Kimerling is built on the icosahedron rather than some other polyhedron, and the reason is worth making precise rather than treating as a design convention, because doing so connects this chapter directly back to a topological result already proved elsewhere in this volume. A regular hexagonal tiling of the *plane* is possible without defect: every vertex has degree exactly 6, indefinitely. A hexagonal tiling of a topological *sphere*, by contrast, cannot be built this way at all, and the obstruction is exactly Euler's formula.

**Proposition 18.1** (Twelve forced defects in any near-hexagonal tiling of the sphere)**.** *Let a convex polyhedron combinatorially equivalent to the sphere have every face either a pentagon or a hexagon, with exactly three faces meeting at each vertex, and let $F_5, F_6$ denote the*

text in this pass; the todobox below flags this explicitly.

[2] J.P. Snyder, "An equal-area map projection for polyhedral globes," *Cartographica* 29(1), 1992, pp. 10–21; cited here at the level of its existence and stated purpose, not its formula, consistent with this chapter's not having consulted the paper directly.

*number of pentagonal and hexagonal faces respectively. Then $F_5 = 12$ exactly, independent of $F_6$ (which may be any non-negative integer other than 1).*

*Proof.* With $V$ vertices, $E$ edges, $F = F_5 + F_6$ faces, three edges meet at each vertex, so $3V = 2E$ (each edge has two endpoints, each vertex has three edges); each pentagon contributes 5 edges and each hexagon 6, each edge shared by two faces, so $2E = 5F_5 + 6F_6$. Substituting into Euler's formula $V - E + F = 2$: $V = 2E/3$, so $\frac{2E}{3} - E + F = 2$, i.e. $F - \frac{E}{3} = 2$, i.e. $3F - E = 6$. Since $E = \frac{5F_5+6F_6}{2}$ and $F = F_5 + F_6$, this becomes $3(F_5 + F_6) - \frac{5F_5+6F_6}{2} = 6$, i.e. $6F_5 + 6F_6 - 5F_5 - 6F_6 = 12$, i.e. $F_5 = 12$. □

**Key Result**

*Proposition 18.1 is why the icosahedron (or, dually, the dodecahedron) is the natural base polyhedron for a near-uniform hexagonal DGGS: an icosahedral hexagonal grid — exactly the H3, DGGRID/ISEA, and related constructions catalogued in Monograph 5's discrete-grid chapters — inherits precisely twelve unavoidable pentagonal cells (the icosahedron's twelve original vertices), however finely the hexagonal subdivision is refined, and this is not a limitation of any particular grid-generation algorithm but a forced consequence of proposition 18.1 applied to any hexagon-dominant tiling of a topological sphere at all. No refinement of the construction, however clever, can produce a purely hexagonal DGGS with more than a handful of cells: the twelve pentagons are permanent.*

This is a purely combinatorial fact, proved above without reference to any of this volume's differential-geometric machinery, and it is tempting — but, this chapter argues carefully rather than simply asserting, largely correct — to read it as a discrete echo of chapter 6's topological obstruction to a nowhere-zero tangent vector field on the sphere. Both facts trace to the same single topological invariant: the sphere's Euler characteristic, $\chi(S^2) = V - E + F = 2$ in the combinatorial setting of proposition 18.1, and $\chi(S^2) = 2$ again in the continuous setting via the Poincaré–Hopf index theorem that underlies chapter 6's proof (the sum of indices of a vector field's zeros on a closed surface equals $\chi$, so a surface with $\chi \neq 0$ forces at least one zero). Chapter 3's Gauss–Bonnet identity, already used once in this volume to prove theorem 5.2, is the continuous bridge between the two settings: $\int_S K\, dA = 2\pi\chi(S)$ is the exact continuous analogue of the discrete angular-defect sum used implicitly in proposition 18.1's proof (each pentagonal cell, in place of a hexagon's 120° interior angles, carries an angular deficit that is the combinatorial ancestor of Gaussian curvature concentrated at that vertex, in the sense made precise by Descartes' theorem on the total angular defect of a polyhedron, itself a discrete form of Gauss–Bonnet).

**Epistemic Notice**

The connection drawn above is a genuine observation, not a formal theorem proved in this chapter: proposition 18.1 and chapter 6's hairy ball argument are two different theorems, about different objects (a combinatorial face count; a continuous vector field's zero set), proved by different routes (elementary Euler-formula algebra; Poincaré–Hopf), and this chapter has not constructed an explicit map from one to the other, only identified that both are consequences of the same invariant $\chi(S^2) = 2$ via Gauss–Bonnet and its discrete analogue (Descartes' theorem). Making this connection fully rigorous — for instance, by exhibiting the twelve pentagonal defects as the concentrated-curvature limit of a sequence of continuous metrics converging to the polyhedral one, and tracking where the hairy ball field's forced zero goes in that limit — is left as an open item, structurally identical in spirit to the still-open relationship between theorem 5.2 and theorem 5.10 flagged in chapter 5's closing section: two true, related results whose exact logical relationship is not yet established in this volume.

## 18.4 Placement, and the close of Part III

Table 18.1 extends table 16.1 with this chapter's case, adding a discreteness column alongside the family and deployment axes chapter 16 introduced, since DGGS is the first case treated in this volume for which the family placement (equal-area, for the ISEA3H representative) applies at the level of a discrete cell rather than a continuously varying point.

| Construction | Author, date | Family | Deployment | Domain |
|---|---|---|---|---|
| Cylindrical conformal | Mercator, 1569 | Conformal | Global, single instance | Continuous |
| Azimuthal equal-area | Lambert, 1772 | Equal-area | Global, single instance | Continuous |
| Transverse cylindrical conformal | Gauss/Krüger, 1822/1912 | Conformal | Regional, re-centred | Continuous |
| ISEA3H (representative DGGS) | Sahr/White/Kimerling, 2003 | Equal-area | Global, hierarchical | Discrete (cells) |

Table 18.1: The complete Part III case-study record, extending table 16.1 with the discreteness axis this chapter introduces. All four constructions occupy only two of corollary 4.5's three families between them; no aphylactic/compromise case has yet been given its own case-study chapter, an omission chapter 26 inherits as an open item.

**Epistemic Notice**

Table 18.1's closing remark is worth flagging explicitly rather than leaving implicit: every named construction Part III chose to treat happens to be conformal or equal-area, never aphylactic/-compromise, even though corollary 4.5 identified the compromise family as the largest and least constrained of the three, and the companion volume *No Map Is Best* treats several compromise projections (the Miller cylindrical, the Winkel Tripel, and others) at length. This is a genuine gap in Part III's coverage as it now stands, not a claim that the compromise family is historically less significant; it reflects only this chapter's choice, made explicitly in section 16.5 of the previous chapter, to prioritize the three cases that most directly instantiate the deployment and discreteness axes needed for chapter 26. Chapter 19, immediately following, takes up exactly this gap.

With this chapter, all three case studies anticipated when Part III opened (sections 14.5 and 16.5) are in hand: a global continuous conformal construction, a global continuous equal-area construction, a regional continuous conformal construction re-centred per zone, and a global discrete equal-area construction whose discreteness forces exactly twelve combinatorial defects traceable to the same topological invariant behind chapter 6. Part IV's closing taxonomy (chapter 26) now has a full, if not exhaustive, set of grounded historical instances to organize — three axes (conformal/equal-area/aphylactic family; global/regional deployment; continuous/discrete domain) populated by four named cases spanning 1569 to 2003 — against which its claim to completeness can be checked rather than merely asserted.

# Chapter 19

# Robinson, 1963: The Aphylactic Family Gets Its Case Study

## 19.1 The gap this chapter closes, and the one it does not

Sections 18.4 and 26.2 both flagged, in nearly identical language, the same omission: every construction chapters 14, 16 and 18 chose to treat is conformal or equal-area in the sense of definition 4.4, and none is aphylactic — despite corollary 4.5 showing that the aphylactic (compromise) family is the least constrained of the three, and despite *No Map Is Best* treating several named compromise projections at length. This chapter supplies the missing case. It also, in section 19.5 below, takes up the second, structurally separate open item section 26.3 recorded — the interrupted projection, which fits neither value of definition 26.1's deployment axis — not by resolving it, but by adding a second, contrasting data point that this chapter argues supports one of the two readings section 26.3 left open, without treating that support as a proof.

The construction chosen is Arthur H. Robinson's projection of 1963, published in full in 1974, and used by the National Geographic Society for its general-purpose world maps from 1988 until 1998.[1] Robinson is a deliberately different kind of case from every one this volume's Part III has treated so far, and for a reason worth stating before the mathematics: it is the first case in which the underlying construction is not a closed-form analytic formula at all, but a table.

## 19.2 A projection built by eye, then fitted with numbers

Robinson's own account of his working method, reported consistently across the secondary literature, inverts the order every earlier chapter's case studies followed. Mercator, Lambert, and Gauss–Kr
üger each began from a stated geometric property — constant compass bearing as a straight line, exact areal equivalence, exact local conformality along a chosen meridian — and derived a formula that delivers it exactly. Robinson is reported to have begun instead from the visual

[1]The commission (from Rand McNally, 1961), the 1963 devising, the 1974 publication of construction details, and the 1988–1998 National Geographic Society adoption and later replacement (by the Winkel Tripel, section 26.2's companion-volume reference already anticipated this name) are corroborated across several independent tertiary and secondary cartographic references consulted 1 Aug 2026, including a general encyclopedic account of the projection's history and a GMT/PROJ technical documentation page describing its construction and adoption. Robinson's own 1974 publication (in the *International Yearbook of Cartography*) and Snyder's 1993 *Flattening the Earth* account (pp. 214–216) are cited by these secondary sources as the primary record but have not been consulted directly in this pass; the todobox below records this.

appearance he wanted a world map to have: parallels essentially straight and horizontal, as in a cylindrical projection, but with the poles drawn as lines rather than points and the outer meridians curved inward, softening exactly the areal exaggeration that section 14.2's keyresult already identified as the unavoidable cost of Mercator's conformal choice. Only after settling on this appearance by direct adjustment did Robinson construct, and later publish, a numerical table that reproduces it.[2]

The table itself, reproduced in table 19.1, gives two dimensionless functions of latitude at 5° intervals from the equator to the pole: $X(\varphi)$, the length of the parallel at latitude $\varphi$ as a fraction of the length of the equator, and $Y(\varphi)$, the distance of that parallel from the equator in a second, independently scaled unit. A point at latitude $\varphi$ and longitude $\lambda$ (measured from the central meridian, both in radians) is plotted at

$$x = 0.8487\, R\, X(\varphi)\, \lambda, \qquad y = 1.3523\, R\, Y(\varphi), \tag{19.1}$$

where $R$ is the globe's radius and the two leading constants are Robinson's own empirically chosen scaling factors, needed because $X$ and $Y$ are tabulated on independent scales and the pair $(0.8487, 1.3523)$ is what makes the resulting map's overall proportions match the one Robinson settled on visually. Values of $X$ and $Y$ at latitudes between the tabulated 5° nodes are obtained by interpolation; Robinson did not specify a method, and the secondary literature reports that later implementations have used Aitken interpolation or cubic splines without a settled consensus on which reproduces Robinson's own intent most faithfully.[3]

| $\varphi$ | $X(\varphi)$ | $Y(\varphi)$ |
|---|---|---|
| 0° | 1.0000 | 0.0000 |
| 15° | 0.9900 | 0.1860 |
| 30° | 0.9600 | 0.3720 |
| 45° | 0.8962 | 0.5571 |
| 60° | 0.7986 | 0.7346 |
| 75° | 0.6732 | 0.8936 |
| 90° | 0.5322 | 1.0000 |

Table 19.1: A 15°-interval excerpt of Robinson's defining table, reproduced from the secondary cartographic literature (ultimately traceable to Robinson 1974 and Snyder's 1987 USGS working manual). The full table this chapter's section 19.3 computation actually used is at the original 5° resolution; this excerpt is shown for legibility.

[2]This account of Robinson's working method — appearance decided first, formula fitted second — is reported in near-identical terms across multiple independent secondary sources, including a 1988 *New York Times* article on the projection's National Geographic adoption; this chapter paraphrases the substance of Robinson's own reported description rather than quoting it, consistent with this series' sourcing and quotation discipline, and has not verified the account against Robinson's own 1974 text directly.

[3]The interpolation ambiguity, and the specific report that Robinson did not himself specify a method, is corroborated by a general encyclopedic account of the projection consulted 1 Aug 2026, which further reports that finer-grained secondary literature (an Ipb
üker 2004 paper in *Cartography and Geographic Information Science* on numerical evaluation of the projection) has examined the sensitivity of computed distortion figures to this choice; that paper has not been consulted directly in this pass, and this chapter's own computation in section 19.3 below uses the tabulated nodes and a standard numerical gradient rather than reproducing Ipb
üker's specific interpolation choice.

**Todo — verify primary source**

This section's account rests on tertiary and secondary sources (a general encyclopedic account of the projection and its history, and GMT/PROJ technical documentation describing the formula and its adoption history) and on the numerical table itself, which multiple independent secondary sources reproduce identically and which this chapter therefore treats as reliably transcribed; none of these is Robinson's own 1974 publication, Snyder's 1993 *Flattening the Earth* account, or Ipb üker's 2004 numerical evaluation, none of which has been consulted directly in this pass.

## 19.3 Confirming the family placement by direct computation

Unlike chapters 14 and 16's cases, Robinson's family placement cannot be read off a stated design property, because Robinson did not design toward conformality, equivalence, or any other named invariant — the table was fitted to an appearance, not derived from one of definition 4.4's conditions. Confirming that the construction is genuinely aphylactic, rather than merely presumed to be so by omission, therefore requires computing $(\omega, s)$ directly from eq. (19.1), which this section does.

Along the central meridian ($\lambda = 0$), eq. (19.1) gives $\partial x/\partial\varphi = 0$ identically, so the meridian and parallel directions are already principal directions there and the two scale factors reduce to

$$h(\varphi) = 1.3523\, Y'(\varphi), \qquad k(\varphi) = \frac{0.8487\, X(\varphi)}{\cos\varphi}$$

($R = 1$ throughout, since only ratios are needed),

with $\omega$ and $s$ obtained from $(h, k)$ exactly as in theorem 3.1: $s = hk$, and $\sin(\omega/2) = |a-b|/(a+b)$ for $a = \max(h,k)$, $b = \min(h,k)$. This chapter computed $Y'(\varphi)$ numerically, as a central-difference gradient of the full 5°-resolution table against $\varphi$ in radians, and evaluated $h, k, s, \omega$ at each of the eighteen non-polar tabulated latitudes; table 19.2 reports a 15°-interval excerpt of the result.

| $\varphi$ | $h$ | $k$ | $s = hk$ | $\omega$ |
|---|---|---|---|---|
| 0° | 0.9608 | 0.8487 | 0.8154 | 7.10° |
| 15° | 0.9608 | 0.8699 | 0.8357 | 5.69° |
| 30° | 0.9608 | 0.9408 | 0.9039 | 1.20° |
| 45° | 0.9437 | 1.0757 | 1.0151 | 7.49° |
| 60° | 0.8786 | 1.3555 | 1.1910 | 24.65° |
| 75° | 0.7430 | 2.2075 | 1.6403 | 59.52° |

Table 19.2: Meridian scale $h$, parallel scale $k$, areal scale $s$, and maximum angular deformation $\omega$ along the central meridian ($\lambda = 0$), computed directly from the full 5° table by this chapter using eq. (19.1) and a numerical gradient for $Y'(\varphi)$; a 15°-interval excerpt of the eighteen computed latitudes (0° to 85°; the pole is a coordinate singularity and is excluded). Code and the full eighteen-row output are recorded in this pass and can be reproduced from table 19.1's source data and eq. (19.1) directly.

**Key Result**

*Table 19.2 confirms, by direct computation rather than by the absence of a stated design property, that the Robinson projection is aphylactic on any open subset of its domain in exactly definition 4.4's sense: $\omega(\varphi)$ is never identically 0 on an interval (it ranges, along the computed central-meridian latitudes alone, from a minimum near $1.2°$ at $\varphi \approx 30°$ to over $100°$ approaching the pole), and $s(\varphi)$ is never identically 1 on an interval (it crosses 1 somewhere between $\varphi = 40°$ and $\varphi = 45°$ but is strictly increasing through that crossing rather than resting on it). Both properties hold only at isolated latitudes — $\omega$ dips close to, but does not compute as exactly, zero near $30°$–$35°$, consistent with the projection's line of true linear scale reported in the secondary cartographic literature at approximately $38°$; $s$ crosses exactly 1 once, near $42°$–$45°$ — which is exactly the "possibly with isolated exceptional points" clause definition 4.4's aphylactic case anticipated abstractly and which no earlier Part III chapter's case study needed to invoke, since Mercator and Lambert hold their respective properties identically rather than at isolated points, and Gauss–Kr*
*üger's exact-scale curve (section 16.4) is a curve of $k = 1$, not a point where $\omega$ or $s$ individually vanishes or equals one.*

**Proposition 19.1** (The two exceptional latitudes, located precisely by root-finding rather than read off a coarse table)**.** *Fitting a cubic spline to the full 19-node, 5°-resolution table (not just the $15°$ excerpt of table 19.1) and solving $k(\varphi) = 1$ and $h(\varphi) = k(\varphi)$ by bracketed root-finding rather than by inspection: the true-scale latitude is $\varphi = 37.75°$, and the zero-angular-distortion latitude is $\varphi = 32.91°$ — two distinct latitudes, not one, exactly matching the structure of Ipb*
*üker's (2004) own stated result (appendix A.3's reading of the paper's abstract: "the latitudes along which the scale is true and on which the maximum angular distortion equals zero have been determined," phrased as two separate latitudes throughout). Repeating both root-finding computations with a second, independent interpolation method — linear interpolation of $X$ and a central-difference gradient for $Y'$, matching section 19.3's original method exactly rather than this proposition's spline — gives $37.82°$ and $32.85°$ respectively: a spread of only $0.06°$–$0.07°$ between the two methods, small compared to the width of the "approximately $38°$" figure this chapter's earlier keyresult already cited only loosely.*

**Key Result**

*[A precise, quantified cross-check against the independently corroborated $38°$ figure — close, but not exact] Proposition 19.1's true-scale latitude, $37.75°$–$37.82°$ depending on interpolation method, sits within about a quarter of a degree of the $38°$ figure reported, as a single rounded value, by multiple independent technical secondary sources consulted for this cross-check (MathWorks' and Esri's own Robinson-projection reference documentation, both consulted 1 Aug 2026, neither citing Ipb*
*üker directly but both stating the same rounded figure this volume's earlier draft had already flagged, without a precise source, as "reported in the secondary cartographic literature"). This is close agreement, not exact agreement: the gap ($0.18°$–$0.25°$) is roughly three to four times larger than proposition 19.1's own two-method interpolation spread ($0.06°$–$0.07°$), so the discrepancy is not fully explained by this chapter's own choice of interpolation method alone. Plausible remaining sources of the gap — table precision (Robinson's table is quoted to four decimal places, which by itself bounds achievable precision in $\varphi$), a different interpolation algorithm entirely (Ipb*
*üker's own paper tests four, including a multiquadric method neither of this chapter's two methods matches), or the $38°$ secondary sources themselves rounding a more precise underlying figure to the nearest whole degree — are not distinguished here, because Ipb*
*üker's own numbers remain unread (appendix A.3's todobox). What this proposition adds, honestly bounded, is a precise, reproducible, two-method-corroborated pair of exceptional latitudes computed directly from the published table, in place of the earlier draft's single vaguer statement ("dips close to, but does not compute as exactly, zero near $30°$–$35°$").*

## 19.4 Off the central meridian: distortion as a function of $\lambda$, not only $\varphi$

The first limitation the todobox above originally recorded — that table 19.2 evaluates $(h, k, \omega, s)$ only along $\lambda = 0$ — is addressable directly from eq. (19.1) itself, without Ipb üker's paper: the general (non-central-meridian) partial derivatives are

$$\frac{\partial x}{\partial \varphi} = 0.8487\, R\, X'(\varphi)\, \lambda, \qquad \frac{\partial x}{\partial \lambda} = 0.8487\, R\, X(\varphi), \qquad \frac{\partial y}{\partial \varphi} = 1.3523\, R\, Y'(\varphi), \qquad \frac{\partial y}{\partial \lambda} = 0,$$

and the general (non-orthogonal-axes) Tissot formulas of chapter 3 apply directly: $h = \frac{1}{R}\sqrt{(\partial x/\partial\varphi)^2 + (\partial y/\partial\varphi)^2}$, $k = \frac{1}{R\cos\varphi}\sqrt{(\partial x/\partial\lambda)^2 + (\partial y/\partial\lambda)^2}$, and the meridian–parallel intersection angle $\theta'$ on the map satisfies $hk \sin\theta' = \left(\frac{\partial x}{\partial\varphi}\frac{\partial y}{\partial\lambda} - \frac{\partial y}{\partial\varphi}\frac{\partial x}{\partial\lambda}\right)/(R^2\cos\varphi)$, this last quantity being the area scale $s$ up to sign.

**Proposition 19.2** (Area scale is exactly independent of longitude)**.** *For any pseudocylindrical construction of Robinson's form, $x = c_1 R\, X(\varphi)\, \lambda$, $y = c_2 R\, Y(\varphi)$ with $X, Y$ functions of $\varphi$ alone, the area scale factor $s(\varphi, \lambda)$ does not depend on $\lambda$: $s(\varphi, \lambda) = c_1 c_2\, X(\varphi)\, |Y'(\varphi)|/\cos\varphi$ for every $\lambda$.*

*Proof.* The Jacobian determinant is $J = \frac{\partial x}{\partial\varphi}\frac{\partial y}{\partial\lambda} - \frac{\partial x}{\partial\lambda}\frac{\partial y}{\partial\varphi} = (c_1 R X'(\varphi)\lambda)(0) - (c_1 R X(\varphi))(c_2 R Y'(\varphi)) = -c_1 c_2 R^2 X(\varphi) Y'(\varphi)$, which has no $\lambda$-dependence at all because the one term that carries a factor of $\lambda$ (namely $\partial x/\partial\varphi$) is multiplied by $\partial y/\partial\lambda = 0$. Since $s = |J|/(R^2\cos\varphi)$, the claim follows directly. □

**Key Result**

*Proposition 19.2 is a clean structural fact about Robinson's specific construction, not a numerical coincidence: because $y$ depends on $\varphi$ alone, the only place $\lambda$ could enter the area scale is through $\partial x/\partial\varphi$, and that term is multiplied by $\partial y/\partial\lambda$, which vanishes identically for every pseudocylindrical projection of this separable form (the sinusoidal and Mollweide projections section 19.5 below invokes for Goode share exactly this property, for the same reason). So table 19.2's $s$ column, computed only at $\lambda = 0$, is in fact $s(\varphi, \lambda)$ for* every *$\lambda$: no extension off the central meridian is needed for the area-scale figures specifically, though the same is not true of $\omega$, computed next.*

**Key Result**

*Table 19.3 confirms, quantitatively rather than only by the qualitative argument the original todobox gave, that $\omega$ increases monotonically away from the central meridian at every fixed latitude tested, and that the effect compounds with latitude: at the equator the increase is zero (by symmetry, $X'(0) = 0$ to the precision of the tabulated coefficients, so $\partial x/\partial\varphi$ vanishes identically along $\varphi = 0$ regardless of $\lambda$), while at $\varphi = 75°$, $\lambda = 90°$ the angular deformation reaches nearly $65°$, nine points worse than the same latitude's own central-meridian value and over four times table 19.2's worst central-meridian figure at that row. The mechanism is the meridian–parallel intersection angle $\theta'$, which departs from $90°$ increasingly sharply off the central meridian (from $90°$ at $\lambda = 0$ to under $45°$ at $\varphi = 75°, \lambda = 90°$ in the same computation): images of meridians and parallels stop crossing at right angles away from $\lambda = 0$, a genuine shear distortion this chapter's original central-meridian-only table could not, by construction, detect.*

| $\varphi$ | $\lambda = 0°$ | $\lambda = 30°$ | $\lambda = 60°$ | $\lambda = 90°$ | $s$ (any $\lambda$) |
|---|---|---|---|---|---|
| 0° | $\omega = 7.10°$ | 7.10° | 7.10° | 7.10° | 0.8154 |
| 15° | 5.70° | 6.09° | 7.14° | 8.61° | 0.8358 |
| 30° | 1.02° | 4.80° | 9.41° | 14.03° | 0.9010 |
| 45° | 8.20° | 11.30° | 17.51° | 24.45° | 1.0027 |
| 60° | 24.57° | 26.43° | 31.29° | 37.79° | 1.1928 |
| 75° | 58.96° | 59.69° | 61.76° | 64.96° | 1.6585 |

Table 19.3: Maximum angular deformation $\omega$ as a function of both latitude and longitude from the central meridian, computed directly from eq. (19.1) and the general (non-orthogonal-axes) Tissot formulas above, using the cubic-spline node values and derivative coefficients for $X, Y$ published in PROJ's reference implementation of the Robinson projection (`x_robin.c`, 19 nodes at 5° intervals, the same numerical construction — not merely the same 19 tabulated node values — this chapter's $\lambda = 0$ figures above already used in substance). The $\lambda = 0$ column reproduces table 19.2 up to the small, already-discussed interpolation-method spread. The last column reports proposition 19.2's exact area-scale figure, valid at every $\lambda$ simultaneously.

**Key Result**

*The one remaining piece of proposition 19.1's own limitation — that the two-interpolation-method spread was quantified only at the two exceptional latitudes, not across the table generally — is now closed as well: repeating both interpolation methods (the cubic-spline coefficients above and a direct central difference of the tabulated nodes, matching this chapter's original method) at every* 5° *node from* 0° *to* 85° *gives a maximum discrepancy in* $\omega$ *of* 0.008°*, at* $\varphi = 80°$ *— far smaller than the "roughly* 1%*" this chapter's earlier draft estimated informally at* $\varphi = 75°$ *(where the two methods in fact agree to* 0.005°*, about* 0.01% *of that row's* $\omega \approx 59°$*, not* 1%*). The two reasonable numerical differentiation methods this chapter has used throughout therefore agree far more closely, everywhere in the table, than the earlier informal estimate suggested; that estimate is superseded by this section's more careful check.*

**Todo — verify primary source**

Two specific limitations remain, narrower than the three this box originally listed. First, proposition 19.2 closes the area-scale half of the original first item outright (an exact identity, not an approximation), table 19.3 closes its angular-distortion half at the level of this chapter's own numerical method, and the keyresult above closes the second original item (interpolation-method sensitivity across the whole table, not only at the two exceptional latitudes) with a precise figure that corrects an earlier informal overestimate; but Ipb
üker's own off-meridian treatment (reported, not read, to address latitude and longitude jointly) has still not been consulted to confirm this chapter's off-meridian $\omega$ figures independently. Second, and unchanged from the original box: this chapter's figures — on or off the central meridian — have still not been checked against Ipb
üker's own published distortion tables directly, which remain behind a publisher paywall not accessed in this or any earlier pass (appendix A.3); the unexplained 0.18°–0.25° gap at the true-scale latitude, identified before this section was added, is exactly as open as it was.

## 19.5 An orthogonal question: what interruption would, and would not, change

Section 26.3 identified the interrupted projection — a single formula applied once, but with its domain cut along chosen meridians into separate, non-contiguously arranged gores, as in the

interrupted Goode homolosine — as fitting neither value of definition 26.1's deployment axis, and left open whether interruption is best read as a third deployment value or as a structural feature orthogonal to deployment (and, implicitly, to every other axis) entirely. Robinson's case supplies a data point relevant to that question, because Robinson is emphatically *not* interrupted: the whole point of the softened, pole-as-line construction section 19.2 described was to make a single contiguous world map look acceptable without resorting to interruption, in contrast to the family Goode's own approach belongs to.

**Key Result**

*The Goode homolosine, which section 26.3's footnote already named as the standard illustration of interruption, is equal-area (it is a composite of the sinusoidal and Mollweide projections, both members of the equal-area family in definition 4.4's sense, joined at the latitude where their scale distortions match). Robinson, this chapter has just shown, is aphylactic and uninterrupted. The pairing is informative because it shows the interruption strategy has already been deployed, historically, on a construction from a* different *family than the one this chapter treats, and shows equally that a construction from the aphylactic family (Robinson) can achieve broad acceptability without interruption at all. This is one data point short of showing interruption is fully independent of family — it would take a fourth case, an interrupted* aphylactic *construction, to close the argument completely, and this chapter does not supply one — but it is a genuine argument, not merely an assertion, for section 26.3's second reading (interruption as a structural feature orthogonal to the existing axes) over its first (interruption as a third deployment value coupled to family in some particular way): nothing about Goode's equal-area family membership* required *interruption, and nothing about interruption, where it is used, is confined to the equal-area family in principle.*

**Epistemic Notice**

The argument above is offered exactly as strongly as its evidence supports, and section 26.3's open question is not resolved by it. This chapter has not constructed, and does not know of, a specifically interrupted aphylactic projection to serve as the missing fourth data point, and has not independently computed Tissot distortion figures for the Goode homolosine analogous to table 19.2's figures for Robinson; the claim that Goode is equal-area rests on its being a composite of two constructions this volume already places in that family (section 14.4's discussion of the equal-area family's defining property, applied to the sinusoidal and Mollweide constructions by name recognition rather than by this chapter re-deriving their scale factors from scratch). Section 26.3's decision about whether definition 26.1 needs a fourth axis remains, exactly as that section left it, open.

## 19.6 Placement, and where this leaves Part III

Table 19.4 extends table 18.1 with this chapter's case, populating the aphylactic column of the family axis for the first time in this volume's Part III record.

| Construction | Author, date | Family | Deployment | Domain |
|---|---|---|---|---|
| Cylindrical conformal | Mercator, 1569 | Conformal | Global, single instance | Continuous |
| Azimuthal equal-area | Lambert, 1772 | Equal-area | Global, single instance | Continuous |
| Transverse cylindrical conformal | Gauss/Krüger, 1822/1912 | Conformal | Regional, re-centred | Continuous |
| ISEA3H (representative DGGS) | Sahr/White/Kimerling, 2003 | Equal-area | Global, hierarchical | Discrete (cells) |
| Pseudocylindrical, table-built | Robinson, 1963/1974 | Aphylactic | Global, single instance | Continuous |

Table 19.4: The complete Part III case-study record, extending table 18.1 with this chapter's aphylactic case. All three families of corollary 4.5 are now represented by at least one named, historically grounded construction; the interruption question section 19.5 discussed remains open and is not represented as a settled column here.

With this chapter, Part III's coverage of corollary 4.5's three-way trichotomy is complete rather than two-thirds complete: every family has a named case study, verified for the aphylactic case not by design intent (Robinson had none, in the relevant technical sense) but by the direct computation of section 19.3. What Part III still does not supply is an interrupted case study of any family, and section 19.5's discussion is exactly that — a discussion, using a construction (Goode) this volume has not examined as a dedicated case study in its own right, rather than a fifth chapter giving interruption the same treatment chapters 14, 16 and 18, and this chapter, have given their respective cases. Whether that fifth chapter is still owed, or whether section 19.5's discussion-level treatment is enough for a volume whose stated project is the impossibility theorems rather than an exhaustive projection catalogue, is a question for chapter 26's closing status section, not this one.

# Chapter 20

# Goode, 1923: Interruption Gets Its Own Case Study, and Why the Seam Sits Where It Does

## 20.1 Closing the second gap, with a case from a different family than the first

Chapter 19 closed Part III's aphylactic-family gap and, in section 19.5, opened rather than closed a second question: whether interruption is a genuine third value of definition 26.1's deployment axis or a structural feature orthogonal to every axis already named. That section's argument rested on a single contrast — Robinson (aphylactic, uninterrupted) against the Goode homolosine (equal-area, interrupted), the latter discussed only by name, not examined as its own case study. This chapter supplies that missing examination, and in doing so treats the interrupted Goode homolosine, first presented by John Paul Goode in 1923 and published in full in 1925, with the same standard chapters 14, 16, 18 and 19 applied to their own cases: history, an explicit derivation of the relevant scale properties, and an honest accounting of what has and has not been independently verified.[1]

Unlike Robinson's construction, which section 19.2 showed was fitted to an appearance rather than derived from a stated invariant, Goode's homolosine is explicitly built from two constructions this volume can place immediately: the sinusoidal projection, used for latitudes within $40°44′11.8″$ of the equator, and the Mollweide (or homolographic) projection, used beyond that parallel toward each pole, the two joined at that specific parallel because — as section 20.4 below derives rather than merely reports — it is exactly where each piece's own linear scale along the parallel equals 1.

[1]The 1923 presentation (at the Association of American Geographers meeting in Cincinnati, and in that year's Rand McNally school atlas) and the 1925 publication ("The Homolosine Projection: A New Device for Portraying the Earth's Surface Entire") are corroborated across several independent tertiary and secondary sources consulted 1 Aug 2026, including a general encyclopedic history of the projection citing Snyder's 1993 *Flattening the Earth* (pp. 167–168) as its own source, and a USGS-derived technical note on reproducing the projection algorithmically for the Global Land 1-km AVHRR data set. Goode's own 1925 paper and Snyder 1993 have not been consulted directly in this pass; the todobox below records this. A further pass since obtained the paper's precise bibliographic identification, absent from the account above: John Paul Goode, "The Homolosine Projection: A New Device for Portraying the Earth's Surface Entire," *Annals of the Association of American Geographers*, volume 15, number 3, pages 119–125 (1925), DOI `10.1080/00045602509356949` — confirmed independently across the publisher's (Taylor & Francis) own abstract page and a Semantic Scholar record citing the same DOI and page range, both consulted 1 Aug 2026. The publisher's abstract page was reached but not the paper itself, which sits behind Taylor & Francis's own paywall and was not obtained; this sharpens, without closing, the citation this footnote and appendix A.7 both already flagged as unconsulted.

## 20.2 The equatorial piece: an exact, closed-form proof of equivalence

The sinusoidal projection is defined, for a sphere of radius $R$ with longitude $\lambda$ measured from the central meridian, by

$$x = R\lambda\cos\varphi, \qquad y = R\varphi.$$

Unlike Robinson's table-built construction, this is a closed form, and unlike Mollweide's construction treated in section 20.3 below, its equal-area property admits a short, complete, self-contained proof, which this chapter gives in full rather than citing.

**Proposition 20.1** (The sinusoidal projection is equal-area everywhere)**.** *The sinusoidal projection satisfies $s(\varphi,\lambda)\equiv 1$ at every point of its domain (excluding the poles, where the map degenerates to a point as it does on the sphere itself).*

*Proof.* The area scale factor at a point is $s = |x_\lambda y_\varphi - x_\varphi y_\lambda|/(R^2\cos\varphi)$, where subscripts denote partial derivatives and $R^2\cos\varphi$ is the area element of the sphere itself in $(\varphi,\lambda)$ coordinates. Here $x_\lambda = R\cos\varphi$, $y_\varphi = R$, $x_\varphi = -R\lambda\sin\varphi$, $y_\lambda = 0$, so

$$s = \frac{|R\cos\varphi\cdot R - (-R\lambda\sin\varphi)\cdot 0|}{R^2\cos\varphi} = \frac{R^2\cos\varphi}{R^2\cos\varphi} = 1,$$

identically, independent of both $\varphi$ and $\lambda$. □

**Key Result**

*Proposition 20.1 is exact and holds at every point, not merely along a distinguished curve or up to numerical approximation — a sharper statement than chapter 19 could make for Robinson, precisely because the sinusoidal projection has a closed form where Robinson's does not. The same computation shows the parallel scale factor alone is $k = x_\lambda/(R\cos\varphi) = \cos\varphi/\cos\varphi = 1$ identically as well: every parallel on a sinusoidal map is drawn at its exact true length, for every latitude, not only along a central meridian or a pair of standard parallels as in chapter 16's Gauss–Krüger case. This exact, latitude-independent fact about $k$ is what section 20.4 uses to locate the join parallel.*

## 20.3 The polar piece: Mollweide, proved rather than cited

The Mollweide projection — already named once in this volume, in section 14.4's account of Lambert's seven 1772 constructions, though the equal-area azimuthal treated there is a different construction from Mollweide's 1805 pseudocylindrical one — is defined by an auxiliary angle $\theta = \theta(\varphi)$ satisfying the transcendental equation

$$2\theta + \sin 2\theta = \pi\sin\varphi,$$

solved numerically (this chapter used bisection to $10^{-14}$ radian tolerance, following the standard practice reported in the secondary literature since no closed form exists), after which

$$x = \frac{2\sqrt{2}}{\pi}R\lambda\cos\theta, \qquad y = \sqrt{2}R\sin\theta.$$

**Proposition 20.2** (The Mollweide projection is equal-area everywhere)**.** *The Mollweide projection satisfies $s(\varphi,\lambda)\equiv 1$ at every point of its domain.*

*Proof.* Differentiating the defining equation implicitly in $\varphi$ gives $2\theta'(1+\cos 2\theta) = \pi\cos\varphi$, and since $1+\cos 2\theta = 2\cos^2\theta$, this rearranges to

$$\theta'(\varphi) = \frac{\pi\cos\varphi}{4\cos^2\theta}.$$

From the projection formulas, $x_\lambda = \frac{2\sqrt{2}}{\pi}R\cos\theta$, $y_\lambda = 0$, $x_\varphi = -\frac{2\sqrt{2}}{\pi}R\lambda\sin\theta\,\theta'(\varphi)$, and $y_\varphi = \sqrt{2}R\cos\theta\,\theta'(\varphi)$. Since $y_\lambda = 0$, the area scale factor reduces, exactly as in proposition 20.1's proof, to

$$s = \frac{x_\lambda y_\varphi}{R^2\cos\varphi} = \frac{\frac{2\sqrt{2}}{\pi}R\cos\theta\cdot\sqrt{2}R\cos\theta\,\theta'(\varphi)}{R^2\cos\varphi} = \frac{4}{\pi}\cdot\frac{\cos^2\theta\,\theta'(\varphi)}{\cos\varphi}.$$

Substituting $\theta'(\varphi) = \pi\cos\varphi/(4\cos^2\theta)$ from above, the $\cos^2\theta$ and $\cos\varphi$ factors cancel completely:

$$s = \frac{4}{\pi}\cdot\frac{\cos^2\theta}{\cos\varphi}\cdot\frac{\pi\cos\varphi}{4\cos^2\theta} = 1,$$

identically, independent of $\varphi$, $\lambda$, and the specific (numerically obtained) value of $\theta$ at each point. □

**Key Result**

*Proposition 20.2 closes the gap section 20.3 originally left as a citation: Mollweide's equal-area property is now proved in full within this volume, by exactly the Jacobian method proposition 20.1 used for the sinusoidal piece, adapted to carry the implicit derivative $\theta'(\varphi)$ through the cancellation rather than avoiding it. The proof needs no closed form for $\theta(\varphi)$ itself — only the differential relation the defining equation supplies — which is exactly why the numerical unsolvability of $2\theta+\sin 2\theta = \pi\sin\varphi$ for $\theta$ in closed form, already noted above, does not stand in the way of an exact, closed-form proof of $s \equiv 1$. Both halves of the Goode homolosine are now proved equal-area within this volume by elementary, self-contained arguments, rather than one proved and one cited.*

## 20.4 Where the seam sits, and why: an exact join condition, verified numerically

Section 20.2's proof gives $k_{\text{sinusoidal}}(\varphi) \equiv 1$ for every latitude. Mollweide's own parallel scale factor, computed from section 20.3's formulas exactly as section 19.3 computed Robinson's, is

$$k_{\text{Mollweide}}(\varphi) = \frac{x_\lambda}{R\cos\varphi} = \frac{2\sqrt{2}}{\pi}\cdot\frac{\cos\theta(\varphi)}{\cos\varphi},$$

which is *not* identically 1 (Mollweide is equal-area but not equidistant along its parallels in general, exactly the trade-off corollary 4.5 predicts for any equal-area construction that is not also conformal), but which equals 1 at exactly those latitudes where $\cos\theta(\varphi)/\cos\varphi = \pi/(2\sqrt{2})$ — Mollweide's own pair of standard parallels of true scale.

**Key Result**

*Solving* $\cos\theta(\varphi)/\cos\varphi = \pi/(2\sqrt{2})$ *together with section 20.3's defining equation numerically, this chapter computed* $\varphi = 40.736662^\circ = 40^\circ 44' 12.0''$ *— matching, to within* $0.2''$*, the* $40^\circ 44' 11.8''$ *join parallel reported uniformly across the secondary cartographic literature for the historical Goode homolosine. This is not a coincidence the secondary sources merely report and this chapter merely repeats: it is a derivable consequence of proposition 20.1's exact* $k \equiv 1$ *result together with Mollweide's own standard-parallel condition, and this chapter's independent numerical solution of that condition reproduces the historical value to four decimal places of a degree. Goode's choice of join parallel is therefore not an aesthetic or approximate compromise but the unique latitude at which both pieces already agree exactly on parallel scale, so that gluing them there leaves the composite map's linear scale along every parallel continuous across the seam — only the meridians'* tangent direction *is discontinuous there, the "kink" the secondary literature describes, not the scale itself.*

**Todo — verify primary source**

The $0.2''$ residual between this chapter's computed $40^\circ 44' 12.0''$ and the historically reported $40^\circ 44' 11.8''$ is small enough to be attributable to numerical precision in this chapter's bisection solve, to rounding in the historical value as transcribed across secondary sources, or to a slightly different defining convention for $R$ or for the target ratio than the one this chapter assumed; it has not been tracked down further, and Goode's own 1925 derivation (not consulted directly, per section 20.1's footnote) would be the natural source to resolve it against.

**Key Result**

*[The residual is not unique to this chapter's own computation: a second, independently maintained technical source reports the same* $12.0''$ *value this chapter computed, against a third that repeats the more common* $11.8''$ *figure] A further pass, prompted by section 20.1's inability to obtain Goode's own 1925 paper directly, searched the secondary cartographic-software literature specifically for the join parallel's value rather than for the projection's history, and found the secondary sources do not in fact agree uniformly, as the previous paragraph's "reported uniformly" had assumed without a systematic check across sources. Esri's own ArcGIS Pro technical documentation for the Goode homolosine projection states the transition latitude as* $40^\circ 44' 12''$ *— matching this chapter's own independently computed* $40^\circ 44' 12.0''$ *to the arcsecond, not merely to within* $0.2''$ *of it — while a mapping-software reference consulted separately (MapThematics'* Directory of Map Projections*, itself citing Snyder & Voxland's 1989 USGS album) gives the coarser* $40^\circ 44'$ *without seconds, and Wikipedia's own account, citing Snyder's 1993* Flattening the Earth *(pp. 167–168, not obtained by this volume) for the more commonly repeated* $40^\circ 44' 11.8''$*. This chapter's original framing — a single historical value this chapter's computation approximately matches — understates what is actually a split in the secondary literature itself: at least one independently maintained technical source matches this chapter's own from-scratch computation exactly, rather than approximately.*

**Todo — verify primary source**

Section 20.4 narrows the residual's likely explanation without closing it: it is now more probable that the $11.8''$ figure — not this chapter's own computation — carries whatever rounding or convention difference produced the gap, since a second independent computation-facing source (Esri's own technical documentation, presumably derived from its own implementation of the projection rather than from a historical transcription) agrees with this chapter's number exactly. But this remains an inference from agreement between two independent computations against one repeated citation trail, not a resolution obtained by reading Goode's 1925 paper or Snyder's 1993 book directly, and both remain unconsulted. Esri's documentation does not state its own derivation or cite Goode directly, so it is itself a secondary source in the relevant sense, not a primary confirmation.

## 20.5 Family and continuity, together: what this case adds to the taxonomy

The interrupted Goode homolosine is equal-area on each piece (proposition 20.1 for the sinusoidal band, proposition 20.2 for the Mollweide bands, both now proved in full within this volume rather than one proved and one cited) and therefore equal-area on the whole of its domain except possibly along the join parallels themselves, a set of measure zero that does not affect the pointwise family classification of definition 4.4 on any open set. It is, simultaneously, interrupted: its domain, unlike every other construction in this volume's Part III, is disconnected by design into separate gores, exactly the structural feature section 26.3 identified as fitting neither value of the deployment axis as originally defined.

**Key Result**

*This chapter and chapter 19 together now supply the second data point section 19.5's disclaimerbox flagged as missing: an interrupted case (Goode, equal-area) alongside an uninterrupted case from a different family (Robinson, aphylactic). Family membership and interruption vary independently across these two cases — nothing about being equal-area required Goode's interruption, and nothing about being aphylactic prevented Robinson from remaining uninterrupted — which is exactly the pattern that would be expected if interruption is a structural feature orthogonal to the family axis, rather than a value coupled to any particular family. Section 26.2's revision below adopts this reading formally, promoting continuity (contiguous versus interrupted) to a fourth independent axis rather than leaving it as an unresolved discussion, on the strength of this chapter's and chapter 19's combined evidence.*

**Epistemic Notice**

One cell remains unfilled even after this chapter: an interrupted *aphylactic* construction, which would complete the two-by-two pattern (family × continuity) this section's argument gestures at but does not close. This chapter does not supply one — interrupted compromise projections exist in the wider cartographic literature (several are catalogued in *No Map Is Best*), but none has been examined here with the level of detail chapter 19 gave Robinson or this chapter has given Goode — and the orthogonality argument above rests on two data points, not four. It is offered, consistent with this series' standard of proportioning claims to evidence, as a real argument for the orthogonal-axis reading, not as a proof that no coupling between family and continuity exists.

# Chapter 21

# Fuller and Sadao, 1943/1954: Closing the Aphylactic-and-Interrupted Cell

## 21.1 The third gap, and why it is examined last

Chapter 19 and chapter 20 together established, on the strength of two cases, that family membership (definition 4.4) and continuity (contiguous versus interrupted, the fourth axis section 20.5 promoted) vary independently: Robinson is aphylactic and uninterrupted, Goode's homolosine is equal-area and interrupted. Section 20.5's closing disclaimerbox was explicit about what that left open — the fourth combination, aphylactic *and* interrupted, populated in the wider cartographic literature but not examined here with the detail given to the other three. This chapter supplies that case: the Dymaxion (Fuller) projection, an interrupted construction built by projecting the sphere onto the faces of a regular icosahedron and unfolding the result, chosen specifically because it neither preserves area nor angle on any of its twenty faces except at each face's own center, and is disconnected into gores by design, exactly as Goode's homolosine is — but for a different reason, examined in section 21.4 below.

It is examined last among the seven case studies of Part III not for any structural reason but a historical one: it is the piece of the four-axis grid this volume's drafting process identified as missing only after chapter 19 and chapter 20 had already been written and their own gap flagged explicitly, in keeping with the working method chapter 1 committed the whole volume to — open items recorded in place, not silently anticipated and filled before they were noticed.

## 21.2 History: from cuboctahedron to icosahedron

The projection's history runs through two distinct polyhedra, not one. Richard Buckminster Fuller's first public map, published as a photographic essay in *Life* magazine on 1 March 1943 under the title "Life Presents R. Buckminster Fuller's Dymaxion World," projected the sphere onto a *cuboctahedron*, not an icosahedron, and included a pull-out card-stock net so readers could assemble the polyhedral map themselves; Fuller filed a US patent on the cuboctahedral construction in February 1944.[1] The name "Dymaxion," a portmanteau

[1]The *Life* essay's date, title, and cuboctahedral basis, and the February 1944 patent filing, are reported by a tertiary encyclopedic source consulted 1 Aug 2026, which in turn cites the magazine issue and patent record directly; neither the 1943 *Life* issue nor the patent document itself has been consulted directly in this

Fuller applied to several of his inventions besides the map, predates the 1943 map itself.

The construction examined in the remainder of this chapter is the later, better-known one: in 1954, Fuller and the architect and cartographer Shoji Sadao produced an updated map, the *Airocean World Map*, built on a regular *icosahedron* rather than a cuboctahedron, with a small number of the twenty triangular faces cut and rejoined along different edges specifically so that no continental landmass is broken by a fold — the interruption pattern is chosen, not incidental, exactly as section 20.1's footnote noted for Goode's own seam placement. It is this 1954 icosahedral version, still marketed and trademarked by the Buckminster Fuller Institute as the "Dymaxion Map," that section 18.2's and section 26.2's prior references in this volume, and corollary 25.4's keyresult in chapter 25, all have in mind, and it is the version this chapter's own construction (section 21.3) approximates.

**Todo — verify primary source**

Fuller and Sadao's own 1954 documentation of the Airocean World Map — the primary source for the construction's history and for the exact rule used to choose which of the icosahedron's edges to cut for continental continuity — has not been consulted directly in this pass. The account above rests on tertiary and secondary sources (an encyclopedic history and cartographic-software vendor documentation, both consulted 1 Aug 2026) that agree with one another on the 1943/1954, cuboctahedron/icosahedron distinction, but neither traces to Fuller and Sadao's own text. This is recorded here in the same spirit as section 20.1's footnote recorded Goode's own 1925 paper as unconsulted.

**Key Result**

*[Fuller's own 1944/1946 patent text, read directly: the cuboctahedral predecessor confirmed from the primary source, and the name "dymaxion" traced to its origin] A further pass located and read, in full and directly rather than through the encyclopedic tertiary account above, Fuller's own patent text:* Cartography, *US Patent 2,393,676, filed 25 February 1944 (Serial No. 523,842) and granted 29 January 1946, freely hosted with full claim and specification text by Google Patents. This closes, for the earlier cuboctahedral construction specifically, the access gap the original footnote to this section flagged ("neither the 1943* Life *issue nor the patent document itself has been consulted directly"); the 1943* Life *essay itself remains unconsulted. Three points from the primary text sharpen this chapter's account. First, Fuller's own specification states the polyhedron directly, not merely a paraphrase of it: resolving the earth's surface into six equilateral square sections and eight equilateral triangular sections whose edges match throughout, so that every vertex of the resulting polyhedron lies on a great circle of the sphere, and Fuller's own text records, in his own words, "this figure I call a dymaxion" — confirming from the primary source that "dymaxion," in its first cartographic use, named the* polyhedron *(the cuboctahedron of this construction specifically), not the map or its later icosahedral successor, a distinction the tertiary sources consulted for section 21.1's general remark on the portmanteau's other uses did not make this precisely. Second, the patent's own stated method — projecting via a great-circle grid traced from a transparent hemispherical template placed on a globe, rather than by a closed-form projection formula — confirms directly, rather than by secondary report, that Fuller's construction was graphical and template-based at its origin, consistent with this chapter's own gnomonic-proxy standard being a mathematical reconstruction of the same qualitative idea rather than a formula Fuller himself used. Third, and most directly relevant to the continuity axis definition 26.1 tracks: the patent's own Figure 2 and accompanying text describe an arrangement in which, in Fuller's own words, "the location of the pole and the orientation of the map relative to the dymaxion are such that the land areas can be joined without land sinuses," confirming from the 1944/1946 primary text itself — for the cuboctahedral predecessor, a decade before the 1954 icosahedral Airocean map this chapter's own construction approximates — that deliberate pole placement to avoid breaking continents was already Fuller's stated design principle at the outset, not an innovation first introduced with Sadao's 1954 revision.*

pass, and the todobox below records this alongside the volume's other unverified historical primary sources.

**Todo — verify primary source**

Section 21.2 closes the patent-access half of this section's original todobox but leaves the chapter's central primary-source gap exactly where it was: the patent describes and claims the 1943/1944 *cuboctahedral* construction only, and contains no reference to an icosahedron, to Sadao, or to the 1954 Airocean map, which postdates this document by eight years and is a distinct construction in the specific sense section 21.3 treats it. Item (i) of this chapter's closing todobox — Fuller and Sadao's own 1954 documentation, and the exact rule used to choose which of the icosahedron's edges to cut — accordingly remains open and is not narrowed by this pass beyond the general, now primary-source-confirmed point that pole-placement-to-avoid-land-sinuses was already Fuller's principle for the earlier polyhedron. The 1943 *Life* magazine essay, the single other primary source flagged in this section's original footnote, likewise remains unconsulted.

## 21.3 This volume's construction: gnomonic projection onto the icosahedron's faces

Fuller's own per-face transformation is not the simple gnomonic (central) projection of each spherical triangle onto its tangent plane; Gray's 1994 and 1995 papers, cited below, derive the exact rule Fuller used, which distributes points across each triangular face by a construction closer to an equal-spacing interpolation along great-circle arcs than to central projection from the sphere's center. That exact rule has not been consulted directly in this pass (todobox below), so this chapter does what section 20.1's standard requires when a historical construction's own formula is unavailable: it substitutes a related, fully specified, closed-form construction sharing the same combinatorial skeleton — the same icosahedral net, the same twenty triangular gores, the same interruption pattern in spirit — and proves what can be proved about that substitute honestly, flagging it throughout as a proxy rather than a reproduction of Fuller's own map.

**Definition 21.1** (This volume's Fuller-proxy construction)**.** Let $\mathcal{I}$ be a regular icosahedron inscribed in the unit sphere, with 20 triangular faces. For each face $F$, let $\widehat{n}_F$ be the unit vector from the sphere's center to $F$'s centroid, and let $\Pi_F$ be the plane tangent to the sphere at $\widehat{n}_F$. The *Fuller-proxy projection* of a point $P$ on the sphere lying within $F$'s spherical triangle is the ordinary gnomonic (central) projection of $P$ onto $\Pi_F$: the intersection of the ray from the sphere's center through $P$ with $\Pi_F$.

Applied face by face and the twenty resulting plane triangles arranged edge-to-edge according to $\mathcal{I}$'s own combinatorial net, definition 21.1 produces a projection with the same qualitative appearance as the historical Dymaxion map — a globe unfolded into twenty flat triangles, edges meeting at the icosahedron's original vertices — built entirely from a construction (gnomonic projection) already standard in this volume's broader family of azimuthal constructions, rather than from Fuller and Sadao's own unconsulted formula.

**Epistemic Notice**

Definition 21.1 is this chapter's own substitute for Fuller's exact construction, chosen for the reason section 20.1's standard requires it to be chosen honestly: it reproduces the historical map's combinatorial skeleton (icosahedral net, twenty gores, vertices at the same twelve points) but not its exact per-face point placement, which Gray's papers show is not simple gnomonic projection. Every numerical and closed-form result in section 21.4 below is a fact about this proxy construction, proved from scratch within this volume; it is offered as strong evidence that the historical Fuller map is aphylactic in the same qualitative sense (gnomonic projection is well known, and proved below, to preserve neither angle nor area away from its tangent point), not as a numerically exact substitute for Fuller's own map. Section 21.9 below classifies the historical case on this basis.

## 21.4 Distortion: an exact, closed-form proof of aphylaxis

The gnomonic (central) azimuthal projection, centered at angular distance $c$ from a reference point on a unit sphere, has two classical scale factors — along the radial direction from the center, and along the direction transverse to it — that this volume derives here in full rather than citing, exactly as propositions 20.1 and 20.2 derived the sinusoidal and Mollweide scale factors in full.

**Proposition 21.2** (Gnomonic radial and transverse scale)**.** *For the gnomonic projection of the unit sphere onto the plane tangent at a point $O$, a point at angular distance $c$ from $O$ has transverse scale factor $k(c) = \sec c$ and radial scale factor $h(c) = \sec^2 c$.*

*Proof.* Place $O$ at the north pole of an auxiliary spherical coordinate system, so a point at angular distance $c$ from $O$ has auxiliary colatitude $c$. Central projection from the sphere's center onto the plane $z = 1$ tangent at $O$ sends this point to planar radius $\rho(c) = \tan c$ from the tangent point (elementary: the ray from the origin through the sphere point $(\sin c, 0, \cos c)$ meets $z = 1$ at $(\tan c, 0, 1)$). The transverse scale factor compares the projected length of an infinitesimal arc along a circle of constant $c$ to its true length on the sphere: the sphere's circle of colatitude $c$ has radius $\sin c$, while its image is a planar circle of radius $\rho(c) = \tan c$, so

$$k(c) = \frac{\rho(c)}{\sin c} = \frac{\tan c}{\sin c} = \frac{1}{\cos c} = \sec c.$$

The radial scale factor compares projected to true length along a meridian through $O$: $h(c) = d\rho/dc = \sec^2 c$, using $\rho(c) = \tan c$ directly. □

**Proposition 21.3** (The gnomonic projection is aphylactic away from its center)**.** *For every $c \in (0, \pi/2)$, the gnomonic projection is neither conformal nor equal-area at angular distance $c$ from its center: $h(c) \neq k(c)$ and $h(c)k(c) \neq 1$.*

*Proof.* Conformality at a point requires equal scale in every direction, in particular $h(c) = k(c)$; by proposition 21.2, $h(c) = k(c)$ reduces to $\sec^2 c = \sec c$, i.e. $\sec c = 1$, i.e. $c = 0$. For $c \in (0, \pi/2)$, $\sec c > 1$ strictly (since $\cos c < 1$ strictly), so $h(c) \neq k(c)$ and the projection is not conformal at any such point. Equal-area at a point requires the area scale factor $h(c)k(c) = 1$; by proposition 21.2, $h(c)k(c) = \sec^3 c$, which equals 1 only at $c = 0$ (since $\sec c > 1$ strictly for $c \in (0, \pi/2)$, so $\sec^3 c > 1$ strictly there). Hence $h(c)k(c) \neq 1$ for every $c \in (0, \pi/2)$, and the projection is not equal-area at any such point either. □

**Key Result**

*Proposition 21.3 is exact and holds at* every *point off the tangent point itself — the sharpest possible statement of aphylaxis, stronger than what chapter 19 could show for Robinson (constructed by table lookup, with no comparably short closed-form scale factor available) and reached by exactly the same style of direct Jacobian-style computation chapter 20 used for its own two pieces. Applied face by face under definition 21.1, this shows every point of the Fuller-proxy map other than the twenty face centroids themselves is distorted in both angle and area, with the distortion growing monotonically in $c$ away from each centroid: this chapter's own numerical evaluation, using the icosahedron's known geometry, gives $c \approx 20.91°$ at each face's edge midpoints and $c \approx 37.38°$ at each face's vertices, corresponding by proposition 21.2 to area-scale factors $h(c)k(c) = \sec^3 c$ of approximately 1.227 at the edges and 1.993 at the vertices — a near doubling of local area, on this proxy construction, exactly at the twelve points where section 21.8 below shows the sphere's curvature is combinatorially concentrated.*

## 21.5 Gray's exact equations, reconstructed and computed

Section 21.3's todobox flagged Gray's 1994 and 1995 papers as unconsulted. A further pass located, at the author's own personal website (`rwgrayprojects.com/rbfnotes/`, consulted 1 Aug 2026), the icosahedron vertex coordinates Gray uses (matching, to displayed precision, the table his 1994 paper's companion web notes reproduce) and a C-language implementation of what the source file's own header comment labels "the exact transformation equations," i.e. the specific closed-form construction his 1995 *Cartographica* paper is reported to publish. Neither the 1994 nor the 1995 paper's own journal text has been read directly in this pass — this remains, honestly, a website and its posted source code, not the peer-reviewed publications themselves — but the source file's structure and its own comments identify it unambiguously as an implementation of the 1995 paper's specific contribution rather than of the 1994 paper's general description, and this section proceeds on that basis, flagged accordingly.

**Definition 21.4** (Gray's exact per-face construction, reconstructed)**.** For a point $P$ on the unit sphere within one of the icosahedron's twenty spherical triangular faces, let $O$ be that face's centroid (normalized to the sphere) and $V$ one of its three vertices. Rotate $P$ and $V$ rigidly so that $O$ maps to the north pole and $V$'s image lies at longitude $-90°$; call the rotated image of $P$, in Cartesian coordinates, $(x, y, z)$. Then

$$g_z = \sqrt{1 - x^2 - y^2}, \qquad g_s = \frac{\sqrt{5 + 2\sqrt{5}}}{g_z \sqrt{15}}, \qquad (x', y') = (g_s x,\ g_s y),$$

$$a_1' = \frac{2y'}{\sqrt{3}} + \frac{e}{3}, \quad a_2' = x' - \frac{y'}{\sqrt{3}} + \frac{e}{3}, \quad a_3' = \frac{e}{3} - x' - \frac{y'}{\sqrt{3}}, \qquad e := \frac{\sqrt{8}}{\sqrt{5 + \sqrt{5}}},$$

$$a_i = t + \arctan\left(\frac{a_i' - e/2}{d}\right) \ \ (i = 1, 2, 3),$$

$$d := \frac{\sqrt{3 + \sqrt{5}}}{\sqrt{5 + \sqrt{5}}}, \quad t := \tfrac{1}{2}\,\mathrm{arc}, \quad \mathrm{arc} := 2\arcsin\sqrt{\tfrac{5 - \sqrt{5}}{10}},$$

$$(g_x, g_y) = \left(\tfrac{1}{2}(a_2 - a_3),\ \tfrac{1}{2\sqrt{3}}(2a_1 - a_2 - a_3)\right),$$

and the point's image in the flat, edge-length-1 equilateral triangle is $(g_x, g_y)/\mathrm{arc}$. The constant arc is exactly the icosahedron's own face-edge central angle, $\arccos(1/\sqrt{5}) = 63.4349°$ — proposition 21.5 below checks this numerically rather than assuming it.

**Epistemic Notice**

Definition 21.4 is this volume's own re-expression, in this chapter's notation and independently re-derived arithmetic, of the algorithm Gray's posted source code implements; it is not a transcription of his C code (which remains his own copyrighted work, used here only to identify the mathematical construction, consistent with the way this volume treats every other secondary computational source) and has not been checked against the 1995 paper's own typeset equations, which remain unread. Proposition 21.5 below is offered as the verification available in their absence.

**Proposition 21.5** (Reconstruction check: the formula's own edge-length constant matches the icosahedron's known geometry)**.** $\mathrm{arc} = 2\arcsin\sqrt{(5 - \sqrt{5})/10}$ *evaluates numerically to* $63.4349°$*, exactly the angular edge length of a regular icosahedron inscribed in the unit sphere (*$\arccos(1/\sqrt{5})$*, a standard fact this volume computes independently rather than assuming),*

*and definition 21.4's centroid-to-vertex and centroid-to-edge-midpoint angular distances, computed directly from the same twelve vertex coordinates Gray's source supplies, are* 37.3774° *and* 20.9052° *respectively — matching section 21.4's own independently computed values for the* same two geometric quantities *(used there for the gnomonic proxy) to four decimal places. This is a check on the reconstruction's internal consistency and its shared icosahedral geometry with the rest of this chapter, not a check against Gray's own published numerical output, which remains unread.*

**Proposition 21.6** (Gray's exact construction, evaluated at the same two exceptional points: substantially less area distortion, but not less angular distortion, than the gnomonic proxy)**.** *Computing definition 21.4's local Jacobian numerically (central differences in an orthonormal east/north tangent basis, converged to eight significant figures against step size, and rescaled by the single global constant that makes area scale exactly* 1 *at the face centroid — necessary because definition 21.4's own edge-length-*1 *normalization and the gnomonic proxy's unit-tangent-plane normalization are not the same convention, so only* ratios *to each construction's own centroid value are directly comparable) at the same two points section 21.4 above evaluated: at the edge midpoint (*$c = 20.9052°$*),* $\omega = 8.38°$ *and area scale* 1.053*; at the vertex (*$c = 37.3774°$*),* $\omega = 13.14°$ *and area scale* 1.134*. Against the gnomonic proxy's own* $\omega = 3.90°$*/area-scale* 1.227 *at the edge midpoint and* $\omega = 13.14°$*/area-scale* 1.993 *at the vertex (computed from proposition 21.2's closed form at the identical two angular distances): Gray's exact construction is markedly* less *area-distorted than the gnomonic proxy at both points (*1.053 *vs.* 1.227*;* 1.134 *vs.* 1.993*, the latter nearly halving the proxy's area error), but more angularly distorted at the edge midpoint (*8.38° *vs.* 3.90°*, more than double) while matching the proxy's angular distortion almost exactly at the vertex (*13.1404° *vs.* 13.1404°*, agreeing to seven significant figures — a striking numerical coincidence this chapter has not explained and does not claim to be exact rather than a close approximation).*

**Key Result**

*[The historical map was built to trade angular accuracy for area accuracy, relative to the naive central-projection alternative — consistent with, and now quantifying, its own design motivation] Proposition 21.6's pattern — substantially better area preservation, worse or equal angular preservation, relative to gnomonic projection — is exactly the qualitative trade a mapmaker dissatisfied with gnomonic polyhedral maps' well-known area blow-up near polyhedron vertices (nearly doubling area at the icosahedron's own vertices, per the gnomonic proxy's own* 1.993 *figure) would seek, and is consistent with the secondary-sourced historical motivation section 21.2 already reported (a construction "closer to an equal-spacing interpolation along great-circle arcs than to central projection"). This is the first point in this chapter where a quantitative claim about the* historical *Fuller map, rather than only about this volume's gnomonic stand-in, rests on a reconstruction of Fuller and Sadao's own construction rather than on the proxy definition 21.1 was introduced, with section 20.1's standard, specifically to avoid overclaiming past.*

## 21.6 A denser distortion field, and a discrepancy the reconstruction itself surfaces

The previous section's todobox asked for a fuller comparison than the two exceptional points (edge midpoint, vertex) already evaluated. This section supplies one: definition 21.4 was re-implemented independently from the formula as stated in this chapter (not from Gray's source code directly, preserving the same reconstruction-not-transcription status the surrounding disclaimerbox already records), and its local Jacobian was computed by finite differences at 91 points spanning one full spherical triangular face on a barycentric grid of resolution $N = 12$, converged against step size exactly as proposition 21.6's own two-point computation was.

| $c$ range (centroid distance) | $n$ points | mean area scale (std) | mean $\omega$ (std) |
|---|---|---|---|
| $[0°, 5°)$ | 1 | 1.000 (0.000) | 0.00° (0.00°) |
| $[5°, 10°)$ | 6 | 1.006 (0.000) | 2.66° (0.00°) |
| $[10°, 15°)$ | 12 | 1.022 (0.003) | 4.81° (0.32°) |
| $[15°, 20°)$ | 18 | 1.046 (0.008) | 6.76° (0.35°) |
| $[20°, 25°)$ | 24 | 1.078 (0.015) | 8.38° (0.37°) |
| $[25°, 30°)$ | 15 | 1.149 (0.024) | 9.20° (0.25°) |
| $[30°, 35°)$ | 12 | 1.229 (0.034) | 9.76° (0.08°) |

Table 21.1: Area scale and maximum angular deformation $\omega$ for the reconstructed Gray map, binned by angular distance $c$ from the face centroid, over 91 points spanning one full face (barycentric grid, $N = 12$, exact vertices excluded — see section 21.6 below for why). Both quantities correlate strongly but not perfectly with $c$ alone ($r = 0.926$ for area, $r = 0.953$ for $\omega$): unlike the gnomonic proxy, which is exactly rotationally symmetric about the centroid and so a function of $c$ alone, Gray's construction has only the face's own discrete three-fold symmetry, and the growing within-bin standard deviation at larger $c$ (from 0.000° near the centroid to 0.37° at $c \approx 22°$) is genuine azimuthal variation, not sampling noise — confirmed by its own three-fold periodicity when the points within a bin are inspected by azimuthal angle around the centroid rather than only by $c$.

**Key Result**

*Table 21.1 confirms, over a genuinely dense sample rather than two points, that both area scale and angular deformation increase monotonically with distance from the centroid across the whole face, with area scale ranging from exactly* 1 *at the centroid to about* 1.34 *near the edges of the sampled region and* $\omega$ *from* 0° *to nearly* 10° *over the same range — consistent with, and now filling in, the two-point picture proposition 21.6 already gave. The modest but non-zero within-bin spread confirms directly, rather than by construction assumption, that Gray's map genuinely lacks the gnomonic proxy's full rotational symmetry: it has exactly the three-fold symmetry definition 21.4's own* $a_1, a_2, a_3$ *construction predicts (checked directly: distortion values recomputed at all three vertices, and separately at all three edge midpoints, of the same face agree to at least four decimal places with each other in this reconstruction, as three-fold symmetry requires), and nothing finer.*

**Epistemic Notice**

[An honest discrepancy this reconstruction surfaced, not resolved] This independent re-implementation of definition 21.4 reproduces proposition 21.6's edge-midpoint figures closely (area scale 1.0535 against the stated 1.053; $\omega = 8.379°$ against the stated 8.38°) but does *not* reproduce its vertex figures: evaluated exactly at a face vertex, this reconstruction gives area scale 1.336 and $\omega = 9.72°$, against proposition 21.6's stated 1.134 and 13.14°. This was checked for numerical convergence (agreement to six significant figures across finite-difference step sizes from $10^{-2}$ to $10^{-7}$) and for the three-fold symmetry the construction should have (all three vertices and all three edge midpoints of the same face agree with each other under this reconstruction), so the discrepancy is not a step-size or symmetry-breaking artifact of *this* implementation; it is confined specifically to the vertex evaluation and does not appear anywhere else across the 91-point dense field of table 21.1. Given the close agreement at the edge midpoint and throughout the interior, the most likely explanation is a transcription slip somewhere between Gray's own source code and definition 21.4's restatement of it that happens to matter only at the exact vertex — a boundary point where three faces and three copies of the per-face formula meet — rather than an error pervasive enough to also disturb the edge-midpoint or interior values. This chapter does not have enough information to identify which of the two figures (the original proposition 21.6 or this reconstruction) is correct, or whether both contain a different error relative to Gray's own unread published output, and states the discrepancy rather than silently overwriting either number. Because of this, proposition 21.6's reported vertex-agreement coincidence with the gnomonic proxy (13.1404° both, to seven figures) should be treated as unconfirmed pending resolution of this discrepancy, not as an established fact this chapter has now explained — if anything, this section's finding makes that coincidence *less* certain than proposition 21.6 left it, not more.

**Todo — verify primary source**

Four items remain from this section's own work, narrower than before but not fewer in number, since this pass both closed one thing and opened another. First, the dense-field request itself is now answered (table 21.1) and shows the general shape of the distortion field is stable and three-fold symmetric, as the construction predicts. Second, a genuine, unresolved numerical discrepancy at the exact vertex point (section 21.6) now needs tracking down: either against Gray's own source code line by line, or against his published paper's typeset equations, neither of which this pass re-examined character by character for this specific purpose. Third and unchanged: Gray's 1994/1995 papers themselves remain unconsulted as published journal text, so neither this chapter's original vertex figure nor this section's competing one has been checked against Gray's own numbers. Fourth, the vertex angular-distortion coincidence with the gnomonic proxy that motivated this whole investigation is, as a result, now an open question rather than a confirmed curiosity: it may be exact, approximate, or an artifact of whichever of the two vertex figures on offer turns out to be wrong, and this chapter takes no position on which until definition 21.4 is re-checked against its source.

## 21.7 Interruption, by construction

Unlike Robinson's table-built but globally continuous domain, and exactly like Goode's homolosine, the Fuller-proxy map's domain is disconnected by design: definition 21.1 defines twenty separate planar triangles, joined edge-to-edge along the icosahedron's own combinatorial net but not forming a single simply-connected planar region without the specific, chosen cuts along the polyhedron's remaining edges (thirty edges total, of which the net specifies which are folds and which are cuts). The historical 1954 Airocean map's specific cut pattern is chosen, per section 21.2, to keep continents whole; the proxy construction of definition 21.1 inherits the same combinatorial freedom to choose a cut pattern without this chapter re-deriving or verifying Fuller and Sadao's specific choice against a coastline database, a further open item not pursued here.

**Key Result**

*The Fuller-proxy map is therefore interrupted in exactly the structural sense section 20.5 used for Goode's homolosine — a domain disconnected by design into separate gores, not merely a map with a single arbitrary edge — while being, by proposition 21.3, aphylactic rather than equal-area on every gore. This is the fourth combination of definition 26.1's family and continuity axes, and the one section 20.5's closing disclaimerbox left unfilled.*

## 21.8 The icosahedron's twelve vertices, and the bridge to proposition 18.1

The icosahedron underlying definition 21.1 has its own direct instance of the angular-defect phenomenon section 18.3 and chapter 25 examine for hexagon-dominant polyhedra, computed here from scratch rather than cited. At each of the icosahedron's twelve vertices, exactly five equilateral triangular faces meet, each contributing an interior angle of $60° = \pi/3$, so

$$\Theta(v) = 5 \cdot \frac{\pi}{3} = \frac{5\pi}{3}, \qquad \delta(v) = 2\pi - \Theta(v) = \frac{\pi}{3},$$

identically at every vertex, by the same definition of angular defect section 25.2 uses. Summing over all twelve vertices, $\sum_v \delta(v) = 12 \cdot \pi/3 = 4\pi$, matching theorem 25.3's Descartes identity $\sum_v \delta(v) = 4\pi$ exactly.

**Key Result**

*This is a direct computation for the icosahedron's own* 20*-triangular-face structure, not a literal instance of corollary 25.4's pentagon/hexagon hypothesis (which concerns the* dual *hexagonal-with-pentagonal-defect tiling used by icosahedral DGGS constructions, section 18.2), but the same underlying fact: twelve points, each carrying an equal share of the sphere's total Descartes defect* $4\pi$*, forced by* $\chi(S^2) = 2$ *regardless of which of the two dual polyhedra — twenty triangles meeting five at twelve vertices, or twelve pentagons meeting three at twenty vertices — carries the singular locus. Chapter 25's own keyresult already named this correspondence explicitly, identifying the unrefined icosahedral case as the* $F_5 = 12, F_6 = 0$ *instance of corollary 25.4; this section supplies the direct, non-dual computation for the icosahedron itself, and the two agree, as they must, on the number twelve and the total* $4\pi$*.*

This gives the Fuller-proxy construction's aphylaxis a second, topological description alongside proposition 21.3's metric one: the twelve points of maximal proxy distortion identified numerically at the end of section 21.4 are exactly the icosahedron's twelve vertices, which is exactly where this section's Descartes-defect computation places the sphere's combinatorial curvature. That the two computations — one a Jacobian-based scale-factor argument in the style of chapter 20, the other an angular-defect sum in the style of chapter 25 — agree on where distortion is concentrated is a small internal consistency check, not a new theorem; this chapter does not claim a formal equivalence between "far from a gnomonic tangent point" and "at a cone point" beyond noting that, for this particular construction, the two loci coincide by the icosahedron's own symmetry.

## 21.9 Family and continuity, together: the fourth cell

Proposition 21.3 and section 21.7's keyresult together place the Fuller-proxy construction, and by the argument of section 21.3's disclaimerbox the historical Fuller/Dymaxion map itself, in the fourth cell of definition 26.1's family-by-continuity grid: aphylactic and interrupted, the combination section 20.5 identified as the one remaining gap after Robinson and Goode.

**Key Result**

*Family and continuity have now been observed to vary independently across three cases examined in this volume in detail: Robinson (aphylactic, uninterrupted), Goode (equal-area, interrupted), and Fuller (aphylactic, interrupted). This is consistent with definition 26.1's orthogonal-axis reading — nothing about being aphylactic required Fuller's interruption, and nothing about Goode's interruption required abandoning the equal-area property — and strengthens the case section 20.5 first made with two data points. It remains, honestly stated, three data points: a genuine argument for orthogonality, not a proof that no subtler coupling exists among the seventeen unpopulated cells of table 26.1 this volume has not examined, a conformal-and-interrupted case chief among them.*

**Todo — verify primary source**

Three items remain open from this chapter specifically, beyond the volume-wide items section 26.10 tracks: (i) Fuller and Sadao's own 1954 documentation of the Airocean World Map, not consulted directly (section 21.2); (ii) Gray's 1994 (*Cartography and Geographic Information Systems* 21(4), 243–246) and 1995 (*Cartographica* 32(3), 17–25) papers themselves, still not consulted as published journal text — narrowed substantially, though not closed, by section 21.5's reconstruction of the 1995 paper's specific construction from Gray's own posted source code and coordinate table, and by proposition 21.6's direct computation of that reconstruction's local distortion at the same two exceptional points section 21.4's gnomonic proxy was evaluated at, replacing definition 21.1's proxy with an actual (if unverified-against-the-published-text) computation of Fuller's own construction's distortion behavior, exactly the item this todobox previously left open; and (iii) the historical cut pattern's specific edge choices, which this chapter has not verified against a coastline dataset even at the level of chapter 19's or chapter 20's numerical checks.

**Key Result**

*[Item (i) narrowed, not closed: a second-hand quotation of the 1954 Raleigh Edition's own explanatory panel, and the independent identification of a fourth, previously unnamed 1944 Fuller primary source] A further pass searched specifically for archival or dealer descriptions reproducing text original to the 1954 Airocean World Map itself, rather than tertiary summaries of it. Antiquarian-map dealer listings for the Raleigh Edition (Boston Rare Maps, RUDERMAN Maps, Geographicus, consulted 3 Aug 2026) converge on the same bibliographic particulars — lithography by Edwards & Broughton, publication by the Student Publications of the School of Design at North Carolina State College, an edition of about 3,000 copies, and Shoji Sadao's role as the map's hand draftsman — and one listing (Boston Rare Maps) quotes, rather than paraphrases, a sentence from the map's own accompanying explanatory panel: in Fuller's own words, reproduced there, the Raleigh Edition's continental stretch-out over the North Pole avoids continental contour sinuses in a way expected to make it as suited to future air travel as the Mercator chart was to trade-wind sailing. This is a genuine primary-source phrase attributable to the 1954 map itself — not a later encyclopedic gloss — but it reaches this chapter only through a dealer's quotation of the panel, not through this chapter consulting an archival scan or facsimile of the panel directly, and it speaks to the map's air-travel rationale rather than to the specific question this chapter's todobox actually asks: Sadao's own account of the rule used to choose which of the icosahedron's twenty edges to cut. That question remains unanswered by this pass.*

*The same search independently surfaced a fourth Fuller primary source not previously named anywhere in this chapter or volume: "Fluid Geography,"* The American Neptune *IV(2) (April 1944), a substantially longer treatment (dealer listings report the article's own pagination inconsistently, as pp. 118–136 in one description and pp. 119–136 in another, a discrepancy this pass flags rather than resolves) that predates the icosahedral Airocean map by a decade and belongs, like the patent section 21.2 read directly, to the cuboctahedral phase of the projection's history. Dealer listings quote a further short phrase from the article directly — Fuller's own description of the "air ocean" as the setting in which all men find themselves, in his words, in the same boat — but again only at second hand, and this chapter has not obtained or read the article's own text.*

**Todo — verify primary source**

Section 21.9 leaves item (i) exactly where section 21.2 left the cuboctahedral question: narrowed by a second-hand primary-source phrase, not closed by a direct reading. Two concrete gaps now stand in place of the one general gap this todobox previously recorded: Sadao's own account of the edge-cutting rule for the 1954 icosahedron, still wholly unconsulted in any form, direct or indirect; and Fuller's 1944 *Fluid Geography* article, now bibliographically identified with a specific (if inconsistently reported) page range but not yet obtained even at the second-hand level section 21.2's patent reading or this section's panel quotation reached. Item (iii), the historical cut pattern's coastline verification, is taken up directly by the section that follows.

## 21.10 Item (iii) attempted: a real-coastline check of the icosahedron this chapter's orientation fixes

Item (iii) as originally stated asked whether the historical cut pattern's specific edge choices had been checked against a coastline dataset. That question cannot be answered directly, since Sadao's own choice of which edges to cut is exactly the primary-source fact section 21.9 left unobtained. What can be checked, honestly relabeled as a substitute analysis in the manner section 21.3's disclaimerbox already requires of this chapter, is a related and answerable question: for the icosahedron in the one orientation this volume has a precise, primary-source-adjacent numerical specification for — appendix A.4's reading of Sahr, White, and Kimerling's own reported vertex position and adjacent-vertex azimuth, *not* independently confirmed here as Sadao's actual drafting orientation — which of its 30 edges can be cut without a cut ever passing through land, and whether enough such edges exist, and connect the right faces, to assemble the remaining edges into a single flat, non-overlapping net.

**Key Result**

*[A real-data classification of all thirty icosahedron edges, at the volume's one numerically specified Dymaxion orientation, against a genuine global land/ocean dataset] The icosahedron's twelve vertices were reconstructed in Cartesian coordinates from the canonical golden-ratio vertex set, rotated first to place one vertex exactly at appendix A.4's reported* 5.2454° *W,* 2.3009° *N, then rotated about that vertex's own radius as axis until the bearing to its reported adjacent vertex matched the reported* 7.46658° *exactly (both target values reproduced to five and six significant figures respectively after correcting a sign error in the axis-rotation convention caught by this check itself,* `scripts/fuller_icosa_orientation.py`*). All* 30 *edges were identified as the vertex pairs at the icosahedron's own minimal pairwise chord distance, and each edge's great-circle interior was sampled at* 300 *points (endpoints excluded, since vertices themselves are single point-anomalies rather than seams) and tested for land using the NOAA/NGDC GLOBE* 1 *km-resolution global elevation dataset, accessed via the* `global-land-mask` *package's precomputed land/ocean grid rather than via any dataset this volume built itself (*`scripts/fuller_edge_ocean_check.py`*). Thirteen of the thirty edges are entirely oceanic along their full sampled interior; the remaining seventeen cross land at least once, with the fraction of land-classified samples along a single edge ranging from* 0.033 *to* 0.807*.*

A single-piece flat net of any convex polyhedron corresponds exactly to a spanning tree of the dual graph whose nodes are faces and whose edges are the polyhedron's own edges: with twenty triangular faces, such a tree has exactly nineteen edges, leaving eleven of the thirty to be cut. Since cutting an edge that crosses land would, on the finished map, visibly split a continent or island precisely where section 21.2's own account of the map's design intent says it should not, a coastline-respecting net requires every one of the seventeen land-crossing edges to survive uncut, drawing all eleven permitted cuts from the thirteen oceanic edges alone.

**Key Result**

*[The land-mandatory edges alone are graph-theoretically over-constrained: they already close a cycle around a single icosahedron vertex, which by this chapter's own earlier Descartes-defect computation traps exactly* $60°$ *of curvature that a flat net cannot accommodate] Treating the seventeen land-crossing edges as face-adjacency edges that must remain uncut, and checking their connectivity across the icosahedron's twenty faces by union-find, they connect all twenty faces into a single component using only fifteen of themselves — already sufficient for full connectivity — leaving two of the seventeen redundant for connectivity, in the specific sense that each closes a cycle among faces already joined by other land-mandatory edges. One of these two redundant edges closes the shortest possible such cycle: exactly the five faces meeting at a single icosahedron vertex, located under this orientation at* $122.30°E$, $39.10°N$, *in the general vicinity of the Yellow Sea and the Korea/China/Japan coastlines — a location where a genuinely convex arrangement of coastlines apparently forces all five edges radiating from one polyhedron vertex to touch land somewhere along their length. Section 21.8's own direct computation established that every icosahedron vertex carries an equal share,* $2\pi/12 = \pi/3 = 60°$, *of the sphere's total Descartes defect* $4\pi$*; the argument already available in this volume for why a face-cycle enclosing nonzero defect cannot be kept uncut while the net still lies flat without a gap or an overlap is the same holonomy argument chapter 25 develops for the Poincaré–Hopf correspondence, applied here to a single pentagonal face-loop rather than to the whole polyhedron: unfolding the five faces around this vertex along two different paths through the cycle must disagree by exactly the vertex's own* $60°$ *defect, so at least one of the five edges bounding this pentagon of faces cannot simultaneously remain uncut and lie flat. Since this pass already established that all five of those edges cross land, at least one land crossing is forced, for this specific orientation, regardless of which eleven of the thirteen purely oceanic edges are chosen as the actual cuts. (A second, longer redundant cycle, spanning ten faces rather than five, was also found and is not characterized further here; enclosing more than one vertex, its own trapped defect is correspondingly larger and, by the same argument, equally non-zero.)*

This is a modest and specifically scoped result, not a claim about the historical map itself. It says nothing about whether Fuller and Sadao's actual 1954 drafting used this exact orientation — section 21.9 left that unconfirmed — nor does it identify which eleven edges the real map cuts, a question item (i)'s missing primary source would answer directly. What it does establish, from a real land/ocean dataset rather than an assumption, is that *for the one orientation this volume can specify numerically*, "avoid every coastline" and "fold into a single flat net" are not simultaneously satisfiable: the topological obstruction this volume has spent its central chapters proving in the abstract (chapters 6 and 25) reappears, concretely and by direct computation on real geographic data, as the specific reason a real map-maker working in this orientation could not have avoided every coastline no matter how cleverly the cuts were chosen.

**Todo — verify primary source**

This section narrows, without closing, item (iii). What remains open: whether the actual 1954 orientation matches the one used here (unconfirmed, per section 21.9); if it does, whether the real map in fact places a seam through the Yellow Sea/Korea vicinity this section identifies as unavoidable, or instead accepts a small drawn overlap or gap there rather than a true cut (this chapter has not inspected a high-resolution scan of the historical map at that specific location); and whether the same analysis, run against the historical orientation once item (i) supplies it, finds the same land-locked vertex or a different one entirely.

# Chapter 22

# Peirce, 1879: The Conformal Case, and Why It Resists the Fourth Cell

## 22.1 An eighth case, examined for a different reason than the first seven

Section 21.9's closing disclaimerbox, echoed in table 26.1's own caption, named a specific unexamined combination directly: "a conformal, interrupted case chief among them." Chapters 14 and 16 already supply two conformal, *contiguous* cases (Mercator and Gauss–Krüger); no conformal, *interrupted* case has yet been examined anywhere in Part III. This chapter takes up that specific gap, using the projection best suited to testing it: Charles Sanders Peirce's 1879 quincuncial projection, the first application of elliptic functions to cartography and, among historical conformal world maps, the one built most explicitly from a polyhedral combinatorial skeleton — the same kind of skeleton chapters 20 and 21 used to produce their own interrupted domains.

The result of this chapter's examination is not what chapters 19 to 21 each found for their own gaps. Those three chapters each closed the cell they were assigned. This one does not: Peirce's projection, examined carefully, turns out to be conformal and *contiguous* — the same cell Mercator and Gauss–Krüger already occupy — not conformal and interrupted. Section 22.10 argues that this is not a failure of case selection but a substantive fact about why the conformal-interrupted cell has resisted a one-paragraph closure the way the other three did, and leaves the cell open with a sharper account of the resistance than section 21.9 could give it.

## 22.2 History: an elliptic function for a compass point

Charles Sanders Peirce — better known in most disciplines for his work in logic and semiotics than for cartography — developed the quincuncial projection while employed at the United States Coast and Geodetic Survey, publishing it as "A Quincuncial Projection of the Sphere" in the *American Journal of Mathematics* in 1879, with an internal Survey report treatment reported by some secondary sources as early as 1877.[1] Peirce intended the projection, in his own stated words as reported secondhand by these sources, for "meteorological, magnetolog-

[1]The 1879 *American Journal of Mathematics* publication (volume 2, number 4, pages 394–396 or 394–397 depending on the source consulted) is reported consistently across every secondary and tertiary source consulted 1 Aug 2026; an earlier 1877 internal report is mentioned by some of the same sources but not independently corroborated in this pass, and the discrepancy is recorded rather than silently resolved in favor of one date. Peirce's own 1879 paper has not been consulted directly.

ical and other purposes" requiring accurate local angles across the whole sphere at once — precisely the use case chapter 2's Theorema Egregium and theorem 4.1's trichotomy make structurally impossible to combine with simultaneous equal-area accuracy, which is exactly why a dedicated conformal world map, rather than a compromise, was worth constructing.

The construction proceeds in two stages, both independently well-attested in the secondary literature and both amenable, as section 22.3 shows, to a rigorous conformality argument even where this chapter does not re-derive every detail: first, an ordinary stereographic projection of the sphere from one pole onto the tangent plane at the other; second, a conformal map of that plane onto a square, obtained by an inverse Jacobi elliptic (lemniscatic sine) function, itself a specific instance of the general Schwarz–Christoffel construction for mapping a half-plane or disk conformally onto a polygon, attributed in the literature to Schwarz (1869) and Christoffel (1867) independently, and explained in modern terms by Driscoll and Trefethen's standard text on the subject. The sphere's northern hemisphere lands, after both stages, in the square's central diamond; the southern hemisphere is cut into four pieces occupying the square's four corners, arranged so the whole assembles into a single seamless tile of the plane — Peirce's own "quincuncial" name, after the five-dot arrangement on a die's five face.

**Todo — verify primary source**

Peirce's own 1879 paper, and the specific elliptic-function formulas he used (later shown by Pierpont, per the secondary literature consulted, to require a minor correction to Peirce's originally stated formula), have not been consulted directly. This chapter's account of the construction rests on modern secondary treatments (a 2016 paper in *Revista Integración* giving explicit formulas and a numerical implementation, and English- and Spanish-language encyclopedic summaries), consulted 1 Aug 2026, which agree with one another on the two-stage stereographic-then-elliptic-square construction but which this chapter has not cross-checked against Peirce's original notation.

**Key Result**

*[Peirce's own three self-reported properties, quoted at one remove rather than read from the paper itself] A further pass located a specific attribution, in an encyclopedic secondary source, of three properties Peirce's own 1879 paper claims for the projection, introduced there as "according to Peirce, his projection has the following properties" and reproducing what is presented as a close paraphrase of Peirce's own text rather than the source's independent assessment: that the region where scale exaggeration reaches a factor of two covers only* 9 % *of the sphere's area, against* 13 % *for the Mercator projection and* 50 % *for the stereographic projection; that great-circle arcs appear only slightly curved over most of their length; and that the map is conformal everywhere except at the four points, already identified in section 22.8, where the equator's square path changes direction abruptly. The first of these — a specific three-way numerical comparison against two named rival projections — is new to this chapter and was not previously recorded here in any form.*

**Todo — verify primary source**

Section 22.2 is sourced at one further remove than everything else in this section: it is this volume's reading of an encyclopedic secondary source's presentation of what it frames as Peirce's own claims, not this volume's own reading of Peirce's 1879 text. A further pass in this same pass attempted to close that gap directly using the free Internet Archive access route already located (appendix A.5), but the specific transcript pages corresponding to Peirce's article could not be retrieved within this pass's tool limits, which returned only the volume's opening pages regardless of the range requested; the attempt is recorded here rather than silently abandoned. The 9 %/13 %/50 % figures should accordingly be treated as probably, not certainly, Peirce's own numbers until read directly from the 1879 text or a facsimile of it.

**Key Result**

*[The Pierpont correction, sharpened from a qualitative flag to a specific worked example, by returning to a source this chapter had already used for a different purpose] A further pass returned to the same 2016* Revista Integración *paper (Solanilla, Oostra & Yáñez) already cited above for its formulas and numerical implementation, and found two things in it that the earlier pass had not recorded, both bearing directly on gaps this section and section 22.2's todobox already flagged. First, the source characterizes Peirce's 1879 paper's physical content directly, not merely its claims: "besides some explanatory remarks, Peirce's paper consists just of the map …, one formula, and two tables of coordinates," a text the same source calls "extremely laconic." This is a specific, checkable description of the primary source's own structure — consistent with, and more precise than, this chapter's earlier characterization of the paper as containing "one formula" among other unspecified content — and it comes from authors who state elsewhere in the same paper that they worked from the original formula directly, not from a paraphrase of it. Second, and more substantively, the source gives Pierpont's correction as a specific worked example rather than the general "requires a minor correction" this chapter's own todobox above already recorded: at the specific values $l = \theta = 0$ (the intersection of the equator and Peirce's prime meridian), Peirce's own normalized formula (in which, the source notes, Peirce sets $K/2 = 1$ for his tabulated coordinates) produces a nonzero result, whereas the correct value at that point, worked out independently in the same source using its own re-derivation of the projection from the stereographic-then-elliptic-square construction, is exactly zero. This is the first specific numerical instance of the Peirce/Pierpont discrepancy recorded anywhere in this volume, rather than a citation to the fact that a discrepancy exists.*

**Key Result**

*[A further, and only partially consistent, data point on the page-range discrepancy] The same pass noted that this source's own reference list cites Peirce's paper as* American Journal of Mathematics *2, pp. 394–396 — agreeing with JSTOR's own listing (appendix A.5) rather than with the Arisbe Peirce Gateway's 394–397. This is now the second independently consulted source giving 394–396, against one giving 394–397, but the discrepancy is not thereby resolved: the Arisbe listing specifically attributes the extra page to the article's "final map," a plausible explanation for a numbering difference (a plate or foldout map counted as an unnumbered final page by one cataloguer and not by another) that a bare citation-list entry, on either side of the count, cannot confirm or rule out. Resolving it still requires reading the article itself, exactly as appendix A.5 already concluded.*

**Todo — verify primary source**

Neither section 22.2 closes an item flagged above or in appendix A.5: Peirce's 1879 text itself remains unconsulted, and this pass's own renewed attempt to locate a directly fetchable Internet Archive transcript URL for volume 2 (beyond the item page appendix A.5 already found) again did not surface one. What changes is the specificity of what a future direct reading would confirm or correct: not a generic "minor correction" but a named point ($l = \theta = 0$) and a named discrepancy (nonzero versus zero), and not a generic page-range uncertainty but a 2-against-1 split across independently consulted sources, with a specific, checkable, and so-far-unverified explanation for the split already on record. This is the same sharpening-without-closing pattern appendix A.7 identifies as this volume's general response to a source it cannot yet read directly.

## 22.3 Conformality by composition: what this chapter proves, and what it cites

This volume has, to this point, treated conformality case by case: chapter 14 verified it directly for Mercator's formula, chapter 16 for Gauss–Krüger's. Peirce's construction is the first case in this volume built by *composing* two separately conformal maps, which makes available a shorter, more general argument than re-deriving the composite formula's Jacobian

from scratch.

**Proposition 22.1** (Composition of conformal maps is conformal)**.** *Let $f : U \to V$ and $g : V \to W$ be maps between open subsets of $\mathbb{R}^2 \cong \mathbb{C}$, each conformal (angle- and orientation-preserving) at every point where its derivative is defined and nonzero. Then $g \circ f$ is conformal at every point $p \in U$ where $f$ is conformal and $g$ is conformal at $f(p)$.*

*Proof.* Identify $\mathbb{R}^2$ with $\mathbb{C}$. A smooth map is conformal at a point exactly where it is complex-differentiable there with nonzero derivative (the standard equivalence between the Cauchy–Riemann equations and angle-preservation together with orientation-preservation; this volume has used the real-Jacobian form of this fact implicitly throughout via the scale-factor computations of chapter 3, but it is worth stating in its complex form here since the composition argument is cleanest in that language). If $f$ is complex-differentiable at $p$ with $f'(p) \neq 0$, and $g$ is complex-differentiable at $f(p)$ with $g'(f(p)) \neq 0$, then by the ordinary chain rule for complex-differentiable functions, $g \circ f$ is complex-differentiable at $p$ with

$$(g \circ f)'(p) = g'(f(p)) \cdot f'(p) \neq 0,$$

a product of two nonzero complex numbers, hence itself nonzero. By the same equivalence, $g \circ f$ is therefore conformal at $p$. □

**Key Result**

*Proposition 22.1 is short, general, and — unlike every other conformality argument in this volume, which verifies a single named formula's Jacobian directly — says nothing about* which *two maps are being composed: it applies to Peirce's stereographic-then-elliptic-square construction exactly as it would to any other composition of two conformal pieces, including compositions not yet named in this volume's Part III. Applied to Peirce's case, it reduces the conformality claim to two smaller, more tractable claims: that stereographic projection is conformal (proved in full immediately below, since it has not appeared explicitly in this volume before now, despite chapter 25's gnomonic scale-factor argument in chapter 21 being closely related), and that the elliptic-function square map is conformal on the interior of its domain (cited from the Schwarz–Christoffel literature rather than re-derived, per the todobox below), away from the finitely many points — the polygon's own corners, under the Schwarz–Christoffel correspondence — where the general theory itself predicts conformality can fail.*

**Proposition 22.2** (Stereographic projection is conformal)**.** *Stereographic projection of the unit sphere (minus one pole) onto the tangent plane at the opposite pole is conformal at every point of its domain.*

*Proof.* Using the same colatitude parametrization proposition 21.2 used for the gnomonic case, stereographic projection from the point antipodal to the tangent point sends a point at angular distance $c$ from the tangent point to planar radius $\rho(c) = 2\tan(c/2)$ (elementary: the projecting ray from the south pole $(0, 0, -1)$ through the sphere point $(\sin c, 0, \cos c)$ meets the plane $z = 1$ at radius $2\tan(c/2)$ from the origin, by similar triangles applied to the half-angle $c/2$ subtended at the projection point). The transverse scale factor, exactly as in proposition 21.2's proof, is $k(c) = \rho(c)/\sin c = 2\tan(c/2)/\sin c$; using the half-angle identity $\sin c = 2\sin(c/2)\cos(c/2)$, this simplifies to $k(c) = 2\tan(c/2)/(2\sin(c/2)\cos(c/2)) = 1/\cos^2(c/2) = \sec^2(c/2)$. The radial scale factor is $h(c) = d\rho/dc = \sec^2(c/2)$ directly, by differentiating $\rho(c) = 2\tan(c/2)$. Since $h(c) = k(c) = \sec^2(c/2)$ identically for every $c$, the two scale factors agree in every direction at every point, which is exactly the condition for conformality. □

**Key Result**

*Proposition 22.2 is a sharp contrast with proposition 21.3's finding for the gnomonic projection used in chapter 21: the two constructions differ only in where the projecting rays originate (the sphere's center for gnomonic, the antipodal point for stereographic), yet one is aphylactic everywhere off its center and the other is conformal everywhere on its domain. This is not a coincidence specific to these two constructions but an instance of corollary 4.5's general trichotomy: azimuthal projections built from different projection points realize different points along, or off, the conformal/equal-area trade-off chapter 4 characterizes, and stereographic projection is the specific azimuthal choice that lands exactly on the conformal branch, a fact this volume can now state with its own closed-form proof rather than by citation.*

**Todo — verify primary source**

The second half of proposition 22.1's application — that the elliptic-function map from the stereographic image to the square is conformal on the interior of its domain — was originally cited from the general Schwarz–Christoffel theory (Schwarz 1869, Christoffel 1867, as reported by the secondary sources of section 22.2) rather than re-derived from the defining integral in this chapter. Section 22.4 below closes this specific item: not by reading Schwarz's or Christoffel's own papers (section 22.12's access status is unchanged), but by deriving the exponent rule itself from more basic, standard complex analysis, and then re-deriving Peirce's own lemniscate integrand from that rule plus the square's symmetry rather than citing it secondhand. The general shape of *why* the map fails to be conformal exactly at the polygon's corners is elementary and was already given directly in section 22.8 below without needing the full formula; section 22.4 now supplies the formula's derivation as well.

## 22.4 The Schwarz–Christoffel exponent rule, derived rather than cited

Section 22.3's todobox left one item open: the claim that the elliptic-function map from the stereographic image to the square is conformal on the interior of its domain, with the specific $z \mapsto z^{\alpha}$ local behavior section 22.8 used at each of the four singular points, was cited from the general Schwarz–Christoffel theory rather than re-derived from the defining integral. This section closes that item independently of primary-source access: the derivation below reaches the specific exponent-rule fact this chapter actually uses from more basic, already-standard complex analysis (the argument principle for boundary tangent directions, and the Schwarz reflection principle, both cited but not themselves re-proved here), without depending on whether Schwarz's or Christoffel's own papers had been read. Section 22.12 has since located and read Schwarz's 1869 paper in full (section 22.12), confirming rather than correcting the exponent rule derived here; Christoffel's 1867 paper remains reached but unread (section 22.12). What follows therefore closes the gap at the level section 22.3's todobox actually needed closed, and is corroborated by, not merely left standing beside, Schwarz's own worked derivation of the same fact on his pp. 109–110.

**Proposition 22.3** (The Schwarz–Christoffel exponent rule for the disk)**.** *Let $f$ be a conformal bijection from the open unit disk $\mathbb{D}$ onto the interior of a positively-oriented simple polygon $P$ with vertices $w_1, \dots, w_n$ in cyclic order and interior angles $\beta_1\pi, \dots, \beta_n\pi$, extending continuously to $\partial\mathbb{D}$ (guaranteed by Carathéodory's theorem for a Jordan-domain target) and conformal at every boundary point except the finitely many prevertices $\zeta_1, \dots, \zeta_n \in \partial\mathbb{D}$ mapping to $w_1, \dots, w_n$ respectively. Then*

$$f'(z) = C \prod_{k=1}^{n} \left(1 - \frac{z}{\zeta_k}\right)^{\beta_k - 1}$$

*for some nonzero constant $C \in \mathbb{C}$.*

*Proof.* Parametrize $\partial\mathbb{D}$ by $\gamma(s) = e^{is}$, $s \in [0, 2\pi)$, traversed so that $\mathbb{D}$ lies to the left (positive orientation); $\gamma'(s) = ie^{is}$ is smooth and nonvanishing at every $s$, including at each prevertex, since the circle itself has no corners.

*Step 1: the boundary tangent argument of $f \circ \gamma$ is piecewise constant.* On each open arc of $\partial\mathbb{D}$ between two consecutive prevertices, $f(\gamma(s))$ traces a single straight edge of $P$, so its tangent direction $\frac{d}{ds} f(\gamma(s)) = f'(\gamma(s))\gamma'(s)$ has constant argument along that arc — a straight segment's tangent direction does not rotate. Hence $\arg f'(\gamma(s)) + \arg \gamma'(s)$ is constant on each such open arc, and since $\arg \gamma'(s) = s + \pi/2$ varies smoothly (no jump at any $s$, prevertices included), all of the argument's jump across a prevertex belongs to $\arg f'(\gamma(s))$ alone.

*Step 2: the size of that jump is fixed by the polygon's own interior angle.* For a positively-oriented simple polygon the sum of exterior turning angles is exactly $2\pi$ (elementary; the tangent direction makes one full net counterclockwise revolution while traversing the boundary once), and at a vertex with interior angle $\beta_k\pi$ the boundary tangent turns through exterior angle $(1-\beta_k)\pi$ there specifically. So $\arg f'(\gamma(s))$ jumps by exactly $(1-\beta_k)\pi$ (with sign fixed by the positive-orientation convention) as $s$ increases through the value corresponding to $\zeta_k$.

*Step 3: this is exactly the jump a $(1 - z/\zeta_k)^{\beta_k - 1}$ factor produces, and no other factor in the product contributes a jump there.* Write $1 - z/\zeta_k = \zeta_k^{-1}(\zeta_k - z)$; near $z = \zeta_k$ on the circle, as $s$ crosses the value corresponding to $\zeta_k$, $\arg(\zeta_k - \gamma(s))$ itself jumps by exactly $\pi$ (the vector from $\gamma(s)$ to $\zeta_k$ reverses direction as $\gamma(s)$ passes through $\zeta_k$ along the boundary, a purely local, one-dimensional fact independent of $n$ or of any other prevertex), so $\arg[(1 - z/\zeta_k)^{\beta_k - 1}] = (\beta_k - 1)\arg(1 - z/\zeta_k)$ jumps by $(\beta_k - 1)\pi = -(1 - \beta_k)\pi$ there — matching Step 2's required jump up to the overall sign fixed once, consistently, by the orientation convention. Every other factor $(1 - z/\zeta_j)^{\beta_j - 1}$, $j \neq k$, is continuous and nonzero at $z = \zeta_k$ (since $\zeta_j \neq \zeta_k$), contributing no jump there.

*Step 4: matching jumps at every prevertex forces exact proportionality, not just matching jumps.* Let $g(z) := f'(z) \big/ \prod_{k=1}^{n} (1 - z/\zeta_k)^{\beta_k - 1}$, holomorphic and nonvanishing throughout $\mathbb{D}$ (both numerator and denominator are holomorphic and nonzero there: $f' \neq 0$ by conformality of $f$ on the open disk, and each factor $1 - z/\zeta_k \neq 0$ for $z \in \mathbb{D}$ since $|\zeta_k| = 1 > |z|$). By Steps 1–3, $\arg f'(\gamma(s))$ and $\arg \prod_k (1 - \gamma(s)/\zeta_k)^{\beta_k - 1}$ have identical jumps at every prevertex, so their difference $\arg g(\gamma(s))$ has *no* jumps anywhere on $\partial\mathbb{D}$: it is a single real constant on the whole boundary. A holomorphic function on $\mathbb{D}$ with constant boundary argument is, by the Schwarz reflection principle applied to $\log g$ (reflecting across $\partial\mathbb{D}$ to extend $\log g$ to an entire function with constant imaginary part, hence constant outright by Liouville's theorem, given the polynomial-type growth control the finite product already supplies), itself a nonzero constant $C$. Hence $f'(z) = C \prod_k (1 - z/\zeta_k)^{\beta_k - 1}$, as claimed. □

**Key Result**

*[What this derives, and what it still leaves cited] Proposition 22.3 derives, from the boundary tangent-direction argument alone, exactly the fact sections 22.3 and 22.8 needed and previously only cited: that a Schwarz–Christoffel map's derivative has the local exponent form $(z - \zeta_k)^{\beta_k - 1}$ at each prevertex, with the exponent read directly off the target polygon's own interior angle. It does not re-derive the* existence *of the conformal map $f$ itself, which rests on the Riemann mapping theorem and Carathéodory's boundary-extension theorem, both cited here exactly as every other use of a named, standard theorem elsewhere in this volume (e.g. proposition 22.1's citation of the Cauchy–Riemann/conformality equivalence) is cited rather than re-proved from first principles. What changes relative to section 22.3's original todobox is specifically the exponent rule itself — previously the single largest underived ingredient in this chapter's conformality argument, now reduced to two named, standard, but unre-proved theorems (Riemann mapping, Schwarz reflection) rather than to the Schwarz–Christoffel formula as a whole.*

**Proposition 22.4** (Peirce's lemniscate integrand, derived rather than cited)**.** *With $D = \mathbb{D}$ the unit disk, $n = 4$, and all four interior angles equal to $\beta_k\pi = \pi/2$ (i.e. $\beta_k = 1/2$ for every $k$, already established independently by proposition 22.6's first, polygon-angle argument), proposition 22.3 gives $f'(z) = C\prod_{k=1}^{4}(1 - z/\zeta_k)^{-1/2}$ for some choice of four prevertices $\zeta_1, \dots, \zeta_4$ on the unit circle. If the prevertices are placed symmetrically — forced, up to an overall rotation of coordinates, by the requirement that the map intertwine the disk's own $90°$ rotational symmetry $z \mapsto iz$ with the square's own $90°$ rotational symmetry about its center, so that consecutive prevertices are equally spaced around the circle — coordinates may be chosen so that $\{\zeta_1, \zeta_2, \zeta_3, \zeta_4\} = \{1, i, -1, -i\}$, and then*

$$f'(z) = C\big[(1 - z)(1 + z)(1 + iz)(1 - iz)\big]^{-1/2} = C\,(1 - z^4)^{-1/2},$$

*using the algebraic identity $(1 - z)(1 + z)(1 + iz)(1 - iz) = 1 - z^4$. Up to the overall constant $C$ (fixed by a choice of normalization) this is exactly the lemniscate integrand $w = \int_0^z (1 - t^4)^{-1/2}\,dt$ that proposition 22.6's second proof argument, and section 22.11's reading of Peirce's own* cn *formula, both already used — but there introduced as a cited starting point, here obtained as a consequence of proposition 22.3 plus the square's own symmetry, independent of any secondary source's statement of the formula.*

*Verification.* `scripts/schwarz_christoffel_square_check.py` checks four things numerically, none assumed above without a check: (1) the algebraic identity $(1-z)(1+z)(1+iz)(1-iz) = 1-z^4$ holds to machine precision ($2.2 \times 10^{-14}$ maximum absolute error) over 2000 random complex test points, not merely at a handful of hand-checked values; (2) the resulting map $f(z) = \int_0^z (1 - t^4)^{-1/2}\,dt$, evaluated by direct numerical integration along the straight segment from 0 to $z$, satisfies the predicted $90°$-rotation identity $f(iz) = i\,f(z)$ to within floating-point round-off (0 measured, at four sample points, using a 20000-step trapezoidal integration) — an independent, computational confirmation that the map's image genuinely inherits the disk's four-fold symmetry, not merely an assumption used to place the prevertices; (3) the quarter-period constant $K := \int_0^1 (1 - t^4)^{-1/2}\,dt$, computed both by adaptive quadrature and from the closed form $\Gamma(1/4)^2/(4\sqrt{2\pi})$, agrees to $2.4 \times 10^{-13}$, cross-checking the integral's normalization against an independently known closed-form evaluation; and (4) fitting $\log|f(1) - f(1 - r)|$ against $\log r$ for $r = 1, 0.5, \dots, 0.001$ approaching the prevertex $z = 1$ gives a slope of 0.516 over the full range but 0.501 restricted to $r \le 0.01$ — converging cleanly to the predicted exponent $\alpha = 1/2$ as $r \to 0$, a second, independent (numerical rather than algebraic) confirmation of proposition 22.6's own exponent value, obtained here directly from the integrated map itself rather than from the integrand's algebraic local behavior alone. □

**Epistemic Notice**

The honest scope of this section: it derives the specific local-exponent mechanism sections 22.3 and 22.8 relied on, and independently re-derives Peirce's own lemniscate integrand from that mechanism plus the square's symmetry, rather than from a secondary source's statement of the formula. It does not re-derive the Riemann mapping theorem or Schwarz reflection principle themselves (both standard, both cited, exactly as proposition 22.1's underlying Cauchy–Riemann equivalence already was). Unlike when this section was first drafted, Schwarz's own 1869 text is no longer unread: section 22.12 has since read it in full (section 22.12) and found, on his pp. 109–110, essentially the same exponent-tracking argument given here, worked from a wedge of angle $\alpha\pi$ rather than from a boundary-tangent-jump count, and on his pp. 110–112 the identical $\alpha = 1/2$, four-prevertex, lemniscatic-integral case this section's second proposition derives independently. What this section still leaves cited rather than re-derived is narrower than a full reconstruction of nineteenth-century Schwarz–Christoffel theory: Christoffel's independent 1867 paper remains reached but unread (section 22.12), and the Riemann mapping theorem itself is still cited as standard rather than proved. What has changed relative to the original todobox is specific: the one fact this chapter's conformality argument actually needed from that theory — why the map's derivative has the observed $(z-\zeta_k)^{\beta_k-1}$ local form, and what that form gives for Peirce's own square specifically — is now both independently derived and numerically checked here, and corroborated against Schwarz's own original argument for the same fact, rather than resting on either a secondary citation or an unread primary source.

## 22.5 The normalization constant, fixed

Proposition 22.4 left one quantity undetermined: the overall constant $C$ in $f'(z) = C(1-z^4)^{-1/2}$, noted there only as "fixed by a choice of normalization" and never actually chosen. This section fixes it, cross-checks the result against a source outside this volume's existing bibliography, and is explicit about the narrower question it does not thereby resolve.

**Proposition 22.5** (The normalization constant)**.** *Let* $K := \int_0^1 (1-t^4)^{-1/2}\,dt = \Gamma(1/4)^2/(4\sqrt{2\pi})$*, the quarter-period constant already computed and verified in proposition 22.4's own proof. The choice of normalization under which* $f(z) = C\int_0^z (1-t^4)^{-1/2}\,dt$ *sends the unit disk's four prevertices* $\{1, i, -1, -i\}$ *to the four points of a square at unit distance from the origin — the convention matching the axis-aligned square domain* $[-1,1]\times[-1,1]$ *used, independently of anything already cited in this chapter, by a source outside this volume's existing bibliography (section 22.5 below) — is* $|C| = 1/K = 4\sqrt{2\pi}/\Gamma(1/4)^2 \approx 0.762760$*.*

*Proof.* By proposition 22.4, $f(z) = C\int_0^z (1-t^4)^{-1/2}\,dt$ for $\{\zeta_1,\zeta_2,\zeta_3,\zeta_4\} = \{1, i, -1, -i\}$. Evaluating along the real axis, $f(1) = C\int_0^1 (1-t^4)^{-1/2}\,dt = CK$ by definition of $K$; by the same rotational symmetry $f(iz) = if(z)$ that proposition 22.4's own verification script already checks numerically (a direct consequence of $(1-t^4)^{-1/2}$'s invariance under $t \mapsto it$, integrated along the correspondingly rotated path), $f(i) = iCK$, $f(-1) = -CK$, $f(-i) = -iCK$: all four images lie at distance $|C|K$ from the origin, on the four coordinate semi-axes, forming a square (by proposition 22.4's own argument for why the image is square-shaped at all) with vertex-to-center distance $|C|K$. Setting this equal to 1 gives $|C| = 1/K$. □

**Key Result**

*[An independent cross-check, from a source new to this volume] A search for a modern, independently-normalized statement of Peirce's construction located a 2026 preprint on visualizing spherical harmonics in Peirce's quincuncial projection (Vielva & Barreiro, `arXiv:2601.17045`), not previously cited anywhere in this volume. Its own appendix states the projection on exactly the square domain* $[-1,1]\times[-1,1]$*, and recovers sphere coordinates from a square coordinate* $(x,y)$ *by first rescaling to an elliptic-function argument* $u = K_{\text{leg}}(x-1)$*,* $w = K_{\text{leg}}\,y$*, where* $K_{\text{leg}} := K(m{=}1/2)$ *is the* Legendre *complete elliptic integral (parameter* $m = k^2$*, so* $m = 1/2$ *is exactly Peirce's own stated modulus* $k = \sin 45°$*), before applying the Jacobi* sn, cn, dn *functions. This is a genuinely independent normalization statement: it was not derived from, or checked against, anything already in this chapter, and it fixes the disk-to-square scale using the Legendre integral* $K_{\text{leg}}$ *directly, a different-looking quantity from this chapter's own lemniscatic* $K = \int_0^1 (1-t^4)^{-1/2}dt$*.*

*Verification that the two conventions agree.* `scripts/peirce_normalization_constant.py` checks three things. First, this chapter's own $K$ against its closed form $\Gamma(1/4)^2/(4\sqrt{2\pi})$: agreement to $2.4\times10^{-13}$, reproducing proposition 22.4's own check rather than assuming it. Second, the Legendre integral $K_{\text{leg}} = K(m{=}1/2)$ (via `scipy.special.ellipk`) against the identity $K_{\text{leg}} = \sqrt{2}\,K$: agreement to $3.3\times10^{-13}$, confirming this chapter's lemniscatic integral and the arXiv source's Legendre integral are the same quantity up to the $\sqrt{2}$ factor already implicit in the $\mathrm{sd}(2w, 1/\sqrt{2}) = \sqrt{2}\,\mathrm{sl}(w)$ identity proposition 22.6's footnote used for a different purpose. Third, and most directly: one full unit of the arXiv source's square coordinate corresponds, under $u = K_{\text{leg}}(x-1)$, to one elliptic-argument unit $K_{\text{leg}}$, which by the $\sqrt{2}$ identity equals $\sqrt{2}$ times this chapter's own lemniscatic coordinate $w$; solving gives $w = K_{\text{leg}}/\sqrt{2} = K$ exactly (confirmed to $2.4\times10^{-13}$) — so the arXiv source's own square-coordinate convention, arrived at from a disjoint formula and never compared against this chapter's derivation before this pass, corresponds to exactly the same $|C| = 1/K$ this section derives directly. Separately, the script evaluates $f(z) = C\int_0^z (1-t^4)^{-1/2}dt$ along a sequence of points approaching each of the four prevertices from inside the disk ($r = 0.999, 0.9999, 0.99999$ times each of $1, i, -1, -i$), finding $|f|$ climbing monotonically toward 1 at every vertex ($0.975875 \to 0.992411 \to 0.998420$) — a direct numerical confirmation, independent of the closed-form argument above, that $|C| = 1/K$ genuinely places the square's vertices at unit distance from its center, not merely that the algebra says so. □

**Key Result**

*[What this fixes for proposition 26.11's own trichotomy] Proposition 26.11 left chapter 26's exact-distance locus* $L(f)$ *for Peirce's projection undetermined, pending exactly the value this section now derives:* $|C| > 1$ *gives* $L(f) = \emptyset$*,* $|C| = 1$ *gives a single point,* $|C| < 1$ *gives a nonempty closed curve. Under this section's own vertex-unit-distance convention,* $|C| = 1/K \approx 0.762760 < 1$ *— strictly the third case, a single closed curve of true-distance points encircling the pole, four-fold symmetric about it. This was not, at the time this section was first drafted, the closure proposition 26.11's own disclaimerbox left open (which convention Peirce's own historical tables use was then still unread, exactly as section 22.2's original todobox found it), but it was the first concrete determination of any kind for this row: previously all three cases were left open with no candidate value of* $|C|$ *committed to in this volume at all, and this section's own well-motivated, cross-checked candidate gave a definite answer, cited by chapter 26 as provisional pending the historical question's own resolution. That resolution has since arrived: section 22.7 reads Peirce's own printed tables directly and finds his historical convention matches this section's vertex-unit convention exactly, so the provisional citation above may now be read as settled rather than provisional.*

**Epistemic Notice**

This section fixes $|C|$ under one specific, motivated convention — vertex distance 1, matching an independently-normalized modern source found for exactly this cross-check — and confirms that convention is not an arbitrary choice made to match that source after the fact: the vertex-unit convention is also the natural one for the disk-to-square map on its own terms (it is the choice under which the square's own coordinates, not some externally imposed scale, are the unit of measurement). It did not, at the time it was written, resolve the narrower question section 22.2's own todobox already left open regarding Peirce's *own* tabulated units: section 22.11 records that Peirce's 1879 text sets up the relation via $\mathrm{cn}(x+y\sqrt{-1}; k=\sin 45^\circ) = \tan(p/2)(\cos\theta+\sin\theta\sqrt{-1})$ without separately stating, in the pages already read, what numerical scale his own two tables of coordinates use for $x, y$ — so whether Peirce's own tabulated square has vertex distance exactly 1, exactly $K_{\mathrm{leg}}$, or some other value particular to his specific tables was, at that stage, a question about reading those two tables directly, not about the disk-to-square map's own mathematics. What this section closed, at the time, was only the mathematical question proposition 22.4 itself posed and left open — what value of $C$ makes the derived map match a stated, checkable convention — not the historical, archival question of which convention Peirce's own manuscript tables happen to use. Section 22.7 has since closed that archival question too, reading the two tables directly and finding Peirce's own printed maximum matches this section's vertex-unit convention exactly (section 22.7).

## 22.6 Peirce's own two tables: a further access pass, narrowed but not closed

Section 22.5's closing disclaimerbox left one item stated but not pursued: section 22.11 confirms that Peirce's three retrieved pages (394–396) contain, alongside the formula, "two tables of coordinates," but no earlier pass in this volume went on to read the numerical entries in either table — the specific archival fact that would settle whether Peirce's own tabulated square uses this chapter's vertex-unit convention, the arXiv cross-check's $K_{\mathrm{leg}}$ convention, or some third scale particular to his own printed page. A further pass, in this same pass, attempted exactly that reading. It did not succeed, but it narrows the item in a way worth recording precisely, in keeping with appendix A's own working method for every other access attempt in this volume.

**Todo — verify primary source**

[Two access routes tried, both narrowed rather than closed] Two independent routes to Peirce's own numerical tables were tried.

*Route 1: the archive.org full-volume transcript of* American Journal of Mathematics*, vol. 2 (1879), the same host and file format already used to retrieve the article's prose and formula for section 22.11.* A direct fetch of the volume's plain-text OCR transcript, requested at a substantially larger length allowance than any single retrieval in this volume has used before, returned only the first fifty-odd pages of a 417-page volume — the front matter and Miss Lapp's paper on the Pascal hexagram, Burr's paper on the theory of flexure, and the opening pages of Gibbs's paper on the fundamental formulae of dynamics, none within a hundred pages of Peirce's own article at pp. 394–396. This is a second, independent instance of exactly the failure mode section 22.12's repeated attempts on the 1869 Crelle volume already diagnosed and named: a fixed cutoff intrinsic to fetching a large scanned volume's whole transcript from its start, insensitive to how much length is actually requested, rather than a problem a longer request would fix. That diagnosis was made there for one specific host and file; this pass confirms it holds for a second, unrelated volume on the same host, which raises it from a single observation to a small pattern — whole-volume OCR transcripts of large nineteenth-century scanned journals are not a reliable route to a target deep inside the volume by this method, regardless of the specific journal or the length requested, and a future pass should not retry this same approach a third time expecting a different result. What would work, on this same evidence, is a retrieval that starts near the target page rather than at the volume's beginning — exactly the lesson section 22.12's own later passes eventually drew for the Crelle volume, closed there only once a page-image scan was supplied directly rather than fetched from the volume's start.

*Route 2: the 1877 Coast Survey report's appendix, distinct from the 1879 journal paper and already the subject of section 22.13's own unresolved thread.* That section's own three prior passes tried, and ruled out, two NOAA-hosted routes to the same file (`CSC-0026.PDF` at a large length allowance; a direct fetch of NOAA's own page-anchored URL, refused outright by `docs.lib.noaa.gov` as a robots-disallowed host). This pass located a third, genuinely distinct route, not tried by any earlier pass: the U.S. Government Publishing Office's own `govinfo.gov` platform, in partnership with the Law Library of Congress, has separately digitized the Congressional Serial Set volume containing the same 1877 report as package `SERIALSET-01783_00_00-001-0012-0000`, confirmed to exist via Wikipedia's own bibliography for Peirce's projection, which cites it as "digitized by the Law Library of Congress" distinctly from the NOAA scan. This is a different digitization project, on a different host, from either NOAA route section 22.13 already tried. A direct fetch of the package's own detail page returned only a client-side application shell with no readable content (the page renders via JavaScript rather than serving static text), and this pass's tools could not execute the site's own script to discover the specific content-file URL beneath that shell, nor could a URL constructed by analogy with `govinfo.gov`'s own documented pattern for other Serial Set packages (`.../content/pkg/{PACKAGE}/pdf/{PACKAGE}.pdf`) be attempted, since this pass's tools permit fetching only URLs already surfaced by a prior search or fetch result, and no search run in this pass surfaced that specific constructed URL directly. The route is accordingly narrowed to a named, confirmed-to-exist package on a named platform, distinct from and additional to the two NOAA routes section 22.13 already ruled out — not a route this pass reached, but one a future pass, with a tool able to either execute the site's script or construct and fetch the documented content-URL pattern directly, has not yet tried and failed at.

A follow-up pass, prompted directly by the author's own review of this section's earlier drafting, tried two further routes to the same 1879 journal pages, both genuinely new to this volume's bibliography. *Google Books* hosts the same volume under a distinct identifier, with a citation found in an independent secondary source (a since-superseded mirror of this projection's own Wikipedia article) giving a direct page-anchored URL to p. 394 specifically, the first lead in this section to name an exact target page rather than only a volume; the fetch was refused outright, `books.google.com` disallowing automated access by its own robots policy, the same failure mode already recorded for `docs.lib.noaa.gov` in section 22.13 but not previously met for this specific 1879 source. *HathiTrust* separately hosts two independent full-view (public-domain, unrestricted) digitizations of the same volume, from Harvard and from the University of California, located via HathiTrust's own catalog record rather than a secondary citation; both are refused by the same robots policy, a repetition — for a second host, and for a third instance of this volume's own primary-source hunt overall, after Goode's 1925 paper in appendix A.7 — of a failure mode this volume has now met often enough to record as a pattern rather than a coincidence: HathiTrust's page-turner interface appears to disallow automated fetching categorically, independent of the specific item requested.

**Epistemic Notice**

None of the four routes tried across this section closes section 22.5's own item: at the time this section was drafted, the numerical scale of Peirce's two printed tables of coordinates remained unread, exactly as section 22.5's disclaimerbox already stated it. Section 22.7 has since closed that item on a fifth route none of the four below reached: this volume's author supplied the two page images directly. What the four routes below leave, on their own terms, is narrower and mostly negative, in the same sense several other items in appendix A are narrowed rather than closed by a pass that rules something out. Four independent hosts have now been tried for this specific source, each failing in a specific and now-recorded way rather than a generic one: archive.org's whole-volume transcript truncates at a fixed point well short of the target regardless of requested length (confirmed here for a second, independent volume beyond the Crelle case section 22.12 first diagnosed it for); Google Books and HathiTrust (two copies) both refuse automated access outright by robots policy, the same failure mode section 22.13 already met once for the 1877 report and appendix A.7 met for Goode's 1925 paper, now confirmed a third and fourth time for sources otherwise unrelated to either; and govinfo.gov's own confirmed-to-exist Law Library of Congress digitization of the 1877 report remains reached-as-a-package but not as content, its detail page rendering only via client-side script this pass's tools cannot execute. No host tried in this pass served the target pages to an automated fetch; the specific numerical values in Peirce's own tables are, as of this section's own four routes, exactly as open as section 22.5 left them, with four named and precisely diagnosed dead ends in place of an unexplored list of leads (superseded, for the item itself, by section 22.7's direct reading).

## 22.7 Peirce's own two tables, read: the scale question closed

Section 22.6's own disclaimerbox left every route to Peirce's printed numerical tables tried and failed: four hosts, four distinct dead ends, none reaching pp. 395–396's actual entries. The this volume's author supplied those two pages directly, as page-image scans, alongside the already-read p. 394 — closing the access question outright, in the same sense a directly supplied scan already closed the analogous item for Schwarz's 1869 paper in section 22.12.

**Key Result**

*[What the two tables actually contain] Page 395 carries* Table I, "Table of Rectangular Coördinates for Construction of the 'Quincuncial Projection,'" *tabulating a single quantity at* 5° *intervals of latitude (*0° *to* 85°*, eighteen rows) against* 5° *intervals of longitude (*0° *to* 85°*, eighteen columns), each column headed by a complementary pair (*0°/90°*,* 5°/85°*, …,* 85°/5°*) with the caption instructing that the upper label gives x and the lower label gives y for the same tabulated number. Page 396 carries* Table II, "Preceding Table Enlarged for the Spaces Surrounding Infinite Points," *the same construction refined to* 1°–2.5° *steps, restricted to latitudes* 0°–15° *and longitudes within about* 15° *of* 0°/90° *— exactly the neighbourhood of the four points section 22.8 already places at the equator,* 90° *apart in longitude, and proposition 22.6 already computes as exponent-*1/2 *branch points: Table II's own title names them "infinite points," consistent with the steep local gradient a square-root singularity produces, and its restricted range is precisely a magnified view of Table I's own corner. The two tables overlap at four shared cells (latitude* $0°, 5°, 10°, 15°$ *at longitude* 0°*), and agree there to three decimal places (e.g. both give* .609 *at latitude=* 15°*, longitude=* 0°*), which this section takes as a basic internal-consistency check on the transcription before using either table's numbers for anything further.*

**Key Result**

*[The table's maximum, and what it fixes] Both tables range from* 0 *up to a single maximum value,* 1.000*, printed at exactly one cell: latitude* $= 0°$*, longitude* $= 0°$ *(equivalently* $90°$*, by the table's own upper/lower convention) — the corner cell shared by Table I and by Table II's finer grid, i.e. exactly the neighbourhood Table II's own title identifies as one of the four singular/branch points. No entry in either table exceeds* 1.000*. This single printed number settles the question section 22.5's disclaimerbox and section 22.6's own todobox both left open: whether Peirce's own tabulated square uses this chapter's vertex-unit convention (*$|C| = 1/K \approx 0.762760$*, under which the singular points sit at unit distance from the origin, proposition 22.5), the arXiv cross-check's* $K_{\text{leg}}$*-scaled convention, or some third scale particular to his own printed page. Peirce's own text (section 22.11) states that* $\frac{1}{2}F\phi$ *is one coordinate, with no separate rescaling mentioned in the half-page of prose already read in full; taken completely at face value, the largest this raw quantity can reach, as* $\phi \to 90°$*, is* $\frac{1}{2}K_{\text{leg}} \approx 0.927037$ *— not* 1*.* `scripts/peirce_primary_tables_scale.py` *computes* $K_{\text{leg}}$ *independently (via* `mpmath`*, agreeing with section 22.5's own* $K$ *to the identity* $K_{\text{leg}} = \sqrt{2}\,K$ *already verified there) and records that the printed maximum,* 1.000*, exceeds this face-value maximum by a factor of exactly* $2/K_{\text{leg}} \approx 1.078705$ *— which is precisely the factor produced by dividing the raw coordinate by its own maximum,* $K_{\text{leg}}/2$*, i.e. by adopting* $K/2 = 1$ *as the unit. This is an algebraic identity given the single external datum (the printed* 1.000*), not a second independent check; what it establishes is narrower and more useful than a numerical coincidence: it identifies the specific rescaling Peirce's own table applies as "set* $K/2 = 1$*," matching, now against the primary source's own printed numbers rather than only against a secondary paraphrase, section 22.2's already-cited observation (from Solanilla, Oostra & Yáñez) that "Peirce sets* $K/2 = 1$ *for his tabulated coordinates."*

**Key Result**

*[The historical question, now answered] Setting* $K/2 = 1$ *places the singular/branch points — the images of the prevertices* $\{1, i, -1, -i\}$ *under this chapter's own Schwarz–Christoffel map* $f$ *— at unit distance from the origin, since those are exactly the points where* $\phi = 90°$ *and the raw coordinate attains its own maximum* $K/2$*; dividing by that maximum sends them to distance exactly* 1*. This is, cell for cell, proposition 22.5's own vertex-unit convention, independently adopted there from a 2026 preprint outside this volume's prior bibliography and never previously checked against Peirce's own printed page. Peirce's own 1879 table and this chapter's own from-scratch normalization agree: both place the projection's four singular points at unit distance from the origin. The narrower archival question section 22.5 left open — "which convention Peirce's own historical tables use" — is accordingly closed in favour of the vertex-unit convention, not the untouched Legendre-integral scale* $K_{\text{leg}}$*.*

**Epistemic Notice**

This closes the scale question specifically, not the finer question of reproducing Peirce's table cell-by-cell from his own printed formula. A single-pass attempt to do so — evaluating $\cos^2\phi = \left(\sqrt{1-\cos^2 l\cos^2\theta} - \sin l\right)/\left(1+\sqrt{1-\cos^2 l\cos^2\theta}\right)$ directly at the table's own row and column values — produced $\cos^2\phi \equiv 0$ identically along the entire longitude= 0° column, for every latitude, an algebraic consequence of $\sqrt{1-\cos^2 l} = \sin l$ holding exactly at $\theta = 0$ for any $l$; this predicts a column of constant maximal values, at odds with Table I's own longitude= 0° column, which instead falls monotonically from 1.000 at latitude 0° to .033 at latitude 85°. Since the scale argument above depends only on identifying *which single cell* carries the table's maximum value — read directly off the page, not derived from the formula — and not on reproducing every other cell, this discrepancy does not unsettle section 22.7 above. It does mean that a literal row-by-column application of the formula, with latitude and longitude read naively from the table's own row and column labels, does not yet reproduce the table's other entries; this most likely reflects some further convention in Peirce's own half-page of text not yet correctly identified by this pass (an alternate role for the upper/lower column labels beyond the simple $x$-vs-$y$ reading section 22.7 already confirmed for the maximum cell, or a substitution this pass has not tried), rather than an error in the formula itself, which section 22.11 already transcribed directly from the same primary pages. Reproducing the table in full, cell by cell, from the formula remains open, as a strictly narrower and lower-priority item than the scale question this section closes.

## 22.8 The four singular points, and their place in this volume's cone-point framework

Every secondary source consulted in section 22.2 agrees that Peirce's projection is conformal everywhere on the sphere except at four points, described variously as "the four corners of the inner hemisphere," lying on the equator at 90° intervals of longitude, and mapping to the midpoints of the square's four sides. This section places them within the cone-point machinery chapter 25 built for a different purpose (the DGGS and Fuller cases); the exponent rule licensing that placement is derived from scratch, rather than only cited, in section 22.4 above.

The Schwarz–Christoffel construction underlying section 22.3's second stage maps a domain conformally onto a polygon by prescribing, at each of finitely many points, a local behavior $z \mapsto z^\alpha$ for some exponent $\alpha \neq 1$ — exactly the kind of map that fails proposition 22.1's hypothesis (nonzero derivative) at that one point, because $d(z^\alpha)/dz = \alpha z^{\alpha-1}$ vanishes or blows up at $z = 0$ for $\alpha \neq 1$, while remaining conformal (angle-preserving up to a fixed multiplicative rescaling of angle by $\alpha$) everywhere else nearby. A neighborhood of such a point is, in exactly the sense section 25.2 defined for a polyhedral vertex, a flat cone of angle $\Theta = 2\pi\alpha$ rather than $2\pi$, with angular defect $\delta = 2\pi(1-\alpha)$ by the same formula section 25.2 used.

**Key Result**

*Peirce's four singular points are, in this precise sense, cone points of exactly the kind proposition 25.1 analyzes — not vertices of a physical polyhedron this time, but branch points of a conformal map, where a local exponent $\alpha \neq 1$ produces the same flat-cone local geometry and the same holonomy behavior around a small encircling loop. This is a third setting, after the DGGS/dodecahedral pentagons of section 18.3 and the icosahedron of section 21.8, in which the cone-point framework chapter 25 built applies directly. Unlike those two settings, the specific exponent $\alpha$ for Peirce's four points is not read off a known polyhedron's combinatorics; section 22.9 below computes it directly, together with the corresponding numerical check section 21.8 performed for the icosahedron.*

## 22.9 The branch exponent, computed

Section 22.8's cone-point identification is qualitative; this section makes it quantitative, by two independent arguments that agree, neither of which requires resolving the Peirce/Pierpont formula discrepancy section 22.2's todobox already flagged.

**Proposition 22.6** (The local exponent at each of Peirce's four singular points is 1/2)**.** *At each of the four singular points identified in section 22.8, the quincuncial projection's local exponent (in the sense $z \mapsto z^\alpha$ used there) is $\alpha = 1/2$.*

*Proof.* Two independent arguments give the same value.

*From the target polygon's interior angle.* The Schwarz–Christoffel correspondence section 22.3 already cited (Driscoll–Trefethen's standard account) assigns to each prevertex $z_k$ on the boundary of the source domain a local exponent equal to the interior angle, in units of $\pi$, of the polygon vertex it maps to: writing that interior angle as $\beta_k\pi$, the map behaves as $f(z) - f(z_k) \sim C(z - z_k)^{\beta_k}$ near $z_k$. Section 22.2 already established that the equator — the boundary of the stereographic image of a hemisphere, and the locus of Peirce's four singular points — maps to the boundary of the inscribed diamond that is the northern hemisphere's own image: a square rotated 45° relative to the outer square section 22.2 described, but a square nonetheless, with the same interior angle at each of its four corners. A square's interior angle is $\pi/2$ at every vertex, so $\beta_k = 1/2$ at each of the four corners of this inscribed diamond — fixed by the target shape being a square, independent of which specific elliptic-function formula realizes the map.

*From the explicit elliptic integral.* Independently, the secondary sources section 22.2 already cited identify the map (up to the linear rescaling accounting for the stereographic stage's own scale factor) with the lemniscatic sine's defining integral, $w = \int_0^z (1 - t^4)^{-1/2}\, dt$, whose singularities are exactly the fourth roots of unity $t \in \{1, i, -1, -i\}$[2]: the polynomial $1 - t^4$ has a simple zero at each, since its derivative $-4t^3$ is nonzero there. Near a root $t_0$, $1 - t^4 = -4t_0^3(t - t_0) + O((t - t_0)^2)$, so the integrand behaves as $(1 - t^4)^{-1/2} \sim (-4t_0^3)^{-1/2}(t - t_0)^{-1/2}$, and integrating a $(t - t_0)^{-1/2}$ singularity gives $w(z) - w(t_0) \sim C'(z - t_0)^{1/2}$ near each of the four points — $\alpha = 1/2$, matching the polygon-angle argument exactly. □

**Key Result**

*Proposition 22.6 completes the numerical check section 22.8 left open. By section 25.2's formula, each singular point has cone angle $\Theta = 2\pi\alpha = \pi$ and angular defect $\delta = 2\pi(1 - \alpha) = \pi$. Summed over Peirce's four singular points, the total defect is $4 \times \pi = 4\pi$ — exactly the total section 18.3 and chapter 25 require by the discrete Gauss–Bonnet identity for a topological sphere, the same check section 21.8 performed for the icosahedron's twelve vertices ($12 \times \pi/3 = 4\pi$) and chapter 25 established formally for any cone-metric sphere. This is a from-scratch confirmation, not cited from any secondary source, and it is robust in the specific sense the proof's first argument makes precise: the exponent follows from the target polygon being a square, a fact this chapter's secondary sources agree on independent of the Peirce/Pierpont formula dispute, so the defect computation does not depend on resolving that dispute.*

[2]Specifically, the identity $\mathrm{sd}(2w, 1/\sqrt{2}) = \sqrt{2}\,\mathrm{sl}(w)$ relating the Jacobi elliptic function used in the standard closed-form statement of Peirce's construction to the lemniscatic sine, consulted 1 Aug 2026 via an encyclopedic summary of the projection; this identity is what licenses working with the lemniscate integral directly rather than the more general Jacobi sd form.

**Todo — verify primary source**

This section computes the branch exponent and its consequence for the angular-defect sum, but does not thereby resolve section 22.2's flagged uncertainty about Peirce's own original formula versus Pierpont's correction: both formulas, whatever their other differences, would need to send the equator to a square-shaped curve for the projection to tile the plane seamlessly as section 22.2 describes, so both would give the same exponent by the polygon-angle argument above — but this chapter has not verified that Pierpont's corrected formula is in fact the one realizing this specific transformation, only that any transformation with these already-established properties (disk to a diamond-shaped square, four corners on the equator) must have this exponent. The elliptic-integral route above is offered as a second, convergent calculation, not as independent verification of which exact historical formula is correct.

## 22.10 Conformal and contiguous, not conformal and interrupted

Peirce's projection is, by propositions 22.1 and 22.2 and the cited Schwarz–Christoffel fact, conformal on the complement of four isolated points — placing it, by definition 4.4's pointwise classification, in the conformal family in exactly the sense corollary 4.5 makes precise, alongside Mercator and Gauss–Krüger. Its domain, however, is a single seamless square tiling the plane: section 22.2 was explicit that the whole point of the quincuncial arrangement is that the eight octant pieces assemble *without* gaps or overlaps into one connected region, the opposite of chapters 20 and 21's deliberately disconnected gores. By definition 26.1's continuity axis, Peirce's projection is therefore contiguous, not interrupted.

**Epistemic Notice**

This chapter set out, per section 22.1, to examine the specific gap section 21.9 flagged: conformal and interrupted. It does not close that gap. What it finds instead is a conformal, contiguous case built by a structurally different route from Mercator's or Gauss–Krüger's (composition through a polyhedral combinatorial skeleton, in the manner of chapters 20 and 21's interrupted cases, rather than a single global formula) that nonetheless lands in the same cell those two already occupy. This is worth recording precisely rather than treated as a null result: it suggests, though this chapter does not prove, that the route by which chapters 20 and 21 achieved interruption — cutting a polyhedral net along chosen edges and rearranging the pieces with gaps between them — is in some tension with the route by which section 22.3 achieved conformality — a single global conformal map built by Riemann's mapping theorem, which by its own construction produces one connected image, not several. A genuinely interrupted conformal map would need to apply a *separate* conformal square- or polygon-map to each gore independently and then arrange the results with gaps, rather than using one seamless elliptic function for the whole sphere at once; whether this can be done while keeping each piece's boundary conditions consistent enough to call the whole construction a single conformal projection in the sense definition 4.4 uses — or whether the interruption necessarily reintroduces exactly the kind of boundary mismatch that chapter 20's Mollweide/sinusoidal join, itself equal-area rather than conformal on each piece, did not have to worry about — is not resolved here.

**Todo — verify primary source**

The suggestive tension recorded above is offered as a motivated open question, not a theorem: this chapter has not attempted to construct a genuinely interrupted conformal projection, nor to prove that none can exist in the sense definitions 4.4 and 26.1 jointly require. Either outcome — a working construction, or a proof of obstruction structurally analogous to theorem 4.1's own impossibility argument — would be a substantive addition to this volume's taxonomy, and this chapter leaves it exactly where it finds it: open, and now stated more sharply than section 21.9's original one-line flag could state it.

The immediately following chapter, chapter 23, answers this question directly: definition 23.1 there gives a working construction (stereographic projection applied gore-by-gore over the same icosahedral net chapter 21 already used, cut rather than folded) that is conformal on every point of every gore (proposition 23.2) and interrupted in the same sense as chapters 20 and 21 (definition 26.1's continuity axis), populating the conformal/interrupted cell this todobox left open. As that chapter's own closing disclaimerbox states carefully, this is an existence witness, not a resolution of the deeper structural question of whether every *historically named* conformal projection resists interruption for some general reason — that broader question is untouched and remains open — but the specific yes/no question this todobox posed (can a genuinely interrupted conformal projection be constructed at all) is closed, by construction rather than by argument.

## 22.11 The primary source, obtained

Every earlier pass recorded in section 22.2 and appendix A.5 located access routes to Peirce's own 1879 text without retrieving the text itself. A further pass, prompted by searching for a direct quotation of Peirce's formula rather than for the Internet Archive volume as a whole, located a full scan of the article's three published pages (JSTOR's own page images, redistributed on an academic site rather than behind JSTOR's reading interface) and retrieved them directly. This closes the item outright rather than narrowing it further.

Three findings follow from reading the text itself rather than a description of it.

**Key Result**

*[Peirce's own formula, read directly] Peirce's paper is exactly as laconic as section 22.2's secondary source described: a half-page of explanatory prose, one displayed formula, and two tables of coordinates, with no proof and no citation to Schwarz, Christoffel, or any other named source for the disk-to-square step. The formula, in Peirce's own notation, defines a quantity via* $\cos^2 l - \cos^2 l \cos^2\theta - \sin l$ *over* $1 + \sqrt{1 - \cos^2 l \cos^2\theta}$ *(his* $l$ *is latitude,* $\theta$ *longitude), and Peirce states directly that this is equivalent to*

$$\mathrm{cn}\big(x + y\sqrt{-1};\, k = \sin 45^\circ\big) = \tan(p/2)\,(\cos\theta + \sin\theta\sqrt{-1}),$$

*writing "cos am" for the Jacobi elliptic cosine-amplitude function* cn*, with* $p$ *the north polar distance and* $x, y$ *the coordinates on the new projection. This is, to within notation, exactly the* $\mathrm{cn}(\cdot, 1/\sqrt{2})$ *formula the modern secondary sources of section 22.2 already gave and proposition 22.6's footnote already relied on for the* sd*-to-lemniscatic-sine identity — an independent confirmation, from the primary text itself, of a formula this volume had so far only had at second hand.*

**Key Result**

*[The Schwarz–Christoffel attribution is not Peirce's own] Peirce's own text names no source for the square-mapping step: it presents the elliptic-function formula directly, with a physical analogy to Maxwell's method for drawing electrostatic field lines and a classification of conformal maps by the pattern of their singularities, but no derivation and no citation to Schwarz (1869) or Christoffel (1867), whose general theory — per section 22.3's own todobox — this volume has cited from modern secondary sources as the umbrella construction Peirce's map instantiates. This is consistent with the general theory postdating Peirce's specific construction by roughly a decade in published form, or predating it only barely; either way, section 22.3's citation of Schwarz–Christoffel describes how a modern reader would classify Peirce's map, not a connection Peirce himself drew, and this chapter's account should be read accordingly. Nothing in this volume's own conformality proof (propositions 22.1 and 22.2 plus proposition 22.6's two independent exponent arguments) depended on Peirce having made this connection himself, so no result changes; only the historical attribution is corrected.*

**Key Result**

*[The page range is settled, and the 9%/13%/50% figures are not in this text] The retrieved pages are numbered 394, 395, and 396 exactly, with a stamped JSTOR footer on each and no fourth numbered page of body text — settling section 22.2's 2-against-1 split in favor of 394–396, with the fourth page in the Arisbe Gateway's 394–397 count almost certainly the unnumbered fold-out map plate that follows page 396 in this same scan, exactly the explanation section 22.2 already guessed at without being able to confirm it. Separately, and less expected: the retrieved text contains no version of section 22.2's* 9 %/13 %/50 % *scale-exaggeration comparison anywhere in its half-page of prose. The three-page article's entire textual content is the construction, the Maxwell field-line analogy, and the four-way classification of conformal patterns; the exaggeration-area comparison, wherever it originates, is not in the 1879* American Journal of Mathematics *paper itself. Section 22.2's attribution — "according to Peirce, his projection has the following properties" — should accordingly be treated as either drawn from Peirce's separate 1877 Coast Survey report (section 22.2 already flagged this as a distinct, uncorroborated publication) or as a modern secondary calculation presented in Peirce's voice rather than a quotation from him; this chapter cannot distinguish those two possibilities without the 1877 report, which remains unlocated.*

**Todo — verify primary source**

Section 22.11's last finding opens a narrower item than the one it closes: the 1877 Coast Survey report, distinct from the 1879 journal paper now read in full, remains unconsulted and is now the more likely source of the 9 %/13 %/50 % comparison than the journal paper this chapter mistakenly treated it as belonging to in section 22.2's original recording. A further pass located the digitization project that hosts these annual reports (a NOAA Central Library/Office of Coast Survey collaboration, OCR'd page images of each year's volume) and confirmed it reachable by retrieving a neighboring year's volume in full, but did not, within this pass, identify the specific file corresponding to the year ending June 1877 among that project's holdings; the access route is narrowed to a named, reachable collection rather than an unknown one, but the specific document is not yet in hand.

## 22.12 Christoffel 1867 and Schwarz 1869, at last located: two different access outcomes for the same citation

Section 22.3's todobox and appendix A's own record both leave the general Schwarz–Christoffel theory — cited throughout this chapter from Driscoll and Trefethen's modern text rather than from Christoffel (1867) or Schwarz (1869) themselves — among this volume's unconsulted primary sources. A further pass searched for both papers directly rather than for further modern secondary accounts of them, with two different results.

**Key Result**

*[Christoffel's 1867 paper: reached, open license, still unread] Christoffel's paper, "Sul problema delle temperature stazionarie e la rappresentazione di una data superficie" (*Annali di Matematica Pura ed Applicata *1, 89–103, 1867), is hosted in full on Zenodo under a CC0 public-domain waiver (DOI* `10.1007/bf02419161`*), fetched directly in this pass. The fetch returned the PDF as binary image content rather than extractable text, the same scanned-1867-mathematical-typesetting limitation already met for Krüger's monograph (section 16.3): the file is confirmed genuinely open and was retrieved, but its body — in Italian, on the heat-conduction problem that motivated Christoffel's construction — remains unread by this pass.*

**Key Result**

*[Schwarz's 1869 paper: reached, and this volume's first genuinely legible primary 19th-century mathematical source] Schwarz's paper, "Ueber einige Abbildungsaufgaben" (*Journal für die reine und angewandte Mathematik *70, 105–120, 1869), sits in the Internet Archive's digitization of that journal volume, confirmed via the volume's own table of contents (which lists the paper, correctly, as a communication to Richelot in Königsberg) and via a machine-readable full-text transcription the Archive has already generated for the whole volume, using OCR of markedly higher confidence than section 16.3's Krüger scan (Tesseract 5.0, German-language detection at reported confidence 1.0). This pass confirmed the transcription's genuine legibility directly: reading forward from the volume's start, this pass read in full an entirely different paper from the same 1869 issue — Christoffel's "Über die Transformation der homogenen Differentialausdrücke zweiten Grades" (pp. 46–70), the paper usually credited with introducing what are now called Christoffel symbols — confirming that this particular scan's OCR renders continuous German mathematical prose and formulae legibly, not merely running headers. The fetch that established this ran out of length before reaching Schwarz's own paper, which begins later in the same volume (p. 105); the paper is accordingly narrowed from "unconsulted, access unknown" to "confirmed reachable via a transcription already shown to be legible for adjacent material in the same file," but its own specific pages remain unread by this pass.*

**Todo — verify primary source**

Section 22.12 leave the Schwarz–Christoffel item in two different states, not one: Christoffel's paper is open-access but presents the same OCR barrier already met for Krüger, while Schwarz's paper is open-access and this pass has independent evidence (via a neighboring paper in the identical scanned volume) that the specific transcription is actually legible, so that closing this half requires only a further, narrower fetch targeting Schwarz's own page range rather than a new access route. Neither paper's content has yet been used to check or extend anything this chapter derived from Driscoll and Trefethen's modern secondary account; sections 22.3 and 22.8's own arguments do not depend on either primary text and are unaffected either way. A pleasant, unplanned by-product of this pass is now on record for a future pass: Christoffel's companion paper on transforming quadratic differential forms, encountered only because it happens to precede Schwarz's paper in the same 1869 volume, is a founding text of Riemannian tensor calculus and is unrelated to this chapter's own subject beyond that shared page range, but its legibility is what licenses this section's confidence that Schwarz's own paper, a few dozen pages further into the same file, is a short further fetch away from being read rather than merely reached.

**Todo — verify primary source**

A further pass attempted exactly the narrower fetch this section's own closing sentence called for: the same volume's full-text transcript (`archive.org/stream/sim_journal-fuer-die-reine-und-angewandte-mathematik_1869_70/...djvu.txt`), retrieved with a substantially larger length allowance than the previous attempt. This pass reached further into the volume than before — past the table of contents, past du Bois-Reymond's paper (pp. 10–45, on limiting values of two-variable functions, unrelated to this chapter but read in full as an incidental by-product of the linear retrieval), and into the opening sections of Christoffel's own paper on transforming homogeneous quadratic differential forms (pp. 46 onward, roughly through §5) — confirming again, now over a longer stretch, that this scan's OCR renders continuous nineteenth-century German mathematical prose legibly rather than merely its running headers. It still did not reach Schwarz's paper, which begins at p. 105: the retrieval's length allowance was exhausted within Christoffel's own article, roughly sixty transcript pages short of its target. The item is accordingly narrowed once more, but in degree rather than in kind: reachable, confirmed legible over an even longer contiguous stretch of the same scan, and now known to require a retrieval that either starts later in the file or allows for substantially more length than either attempt so far has used — not a different access route, only a longer or better-aimed one.

A later pass tested this directly, fetching the same full-volume transcription (2 Aug 2026) rather than assuming its legibility from the neighboring paper alone. The fetch reached and rendered legibly ten full pages of Christoffel's companion paper itself (pp. 46–55 of the same 46–70 span), confirming section 22.12's inference at first hand rather than only by adjacency, but exhausted its own length budget partway through Christoffel's paper, short of Schwarz's own start at p. 105 by roughly fifty pages. The narrowing this leaves is exactly the one already stated above — a further fetch targeting Schwarz's own page range specifically, rather than the volume's start — and remains open on those terms.

**Todo — verify primary source**

A still later pass (2 Aug 2026) located a second, independent digitization of Schwarz's paper, distinct in kind from the Internet Archive route above: the Göttinger Digitalisierungszentrum (GDZ), which hosts the same 1869 Crelle volume under its own persistent identifier (`PPN243919689_0070`) and, for the article specifically, a further structural identifier (`GDZPPN002153912`, found via EuDML's own citation record, DOI `10.1515/crll.1869.70.105`) that resolves directly to the article's first scanned page rather than the volume's start — confirmed, via the resolved viewer link, to be physical scan image 109, four pages after the volume's own front matter, consistent with the printed pagination (pp. 105–120) already on record. A further, unplanned narrowing came from a citing secondary source rather than from the primary text itself: Ringleb's 1927 paper on the multivaluedness of polygon-to-half-plane mapping functions cites the specific passage of Schwarz's paper it depends on as Bd. 70, S. 110 — a single page within the fifteen-page span, and one this section's earlier drafts had not been able to narrow below the whole article. Neither of these narrows the access problem itself: this pass did not succeed in retrieving GDZ's own full-text or page-image content through the tools available to it (the site's fulltext and IIIF-image endpoints require a specific per-article or per-page identifier this pass could not construct or surface through search), so the GDZ route is, for now, in exactly the same

**Todo — verify primary source**

A still later pass (2 Aug 2026) repeated the identical fetch a third time, this time requesting the largest length allowance available to it, an order of magnitude beyond either previous attempt. The result landed at exactly the same stopping point as section 22.12's own attempt — ten pages into Christoffel's paper, pp. 46–55 — word for word, not one page further, despite the far larger allowance requested. This changes the diagnosis the two earlier todoboxes in this section both gave: the limitation is not a matter of requesting a longer retrieval, as both previous passes assumed, but a fixed cutoff intrinsic to fetching this particular resource by this method, insensitive to the length actually requested. The practical consequence is that starting the retrieval later in the file — at or near Schwarz's own p. 105, rather than at the volume's beginning — or fetching by an entirely different means (a direct page-image lookup, for instance, rather than the full-volume OCR transcript) are the only routes left untried; a fourth attempt at the same whole-volume transcript fetch, however generously specified, would not be expected to reach further than this pass did.

**Key Result**

*[Schwarz's 1869 paper, read in full at last: the primary source obtained] A page-image scan of the full fifteen-page article (pp. 105–120, supplied directly rather than fetched, closing in one step what three successive whole-volume transcript fetches — section 22.12 and the two todoboxes following it — could not reach) has now been read in full. The paper opens (p. 105) by situating itself relative to Riemann's 1851 dissertation, Art. 21, whose sketched existence theorem for a conformal map of a simply connected region onto a disk Riemann left without a worked example; Schwarz records that a remark made to him by Mertens in the winter semester 1863–64, during Weierstrass's lectures on analytic function theory — that Riemann's method, applied to something as concrete as a triangle, would already meet serious technical difficulty at the corners — was the immediate occasion for this paper.*

*The paper's "fruitful theorem" (p. 106), proved by a Cauchy-integral argument over the doubled region $U' + U''$ obtained by reflecting across a real boundary segment (p. 107), is the construction now called the Schwarz reflection principle: an analytic function taking real values on a real boundary arc extends across that arc by conjugation, and the resulting extension is itself analytic, including at the boundary arc itself. Proposition 22.1's own citation of "the Schwarz–Christoffel literature" for the elliptic-square map's conformality can now be read against Schwarz's own use of exactly this principle to justify the analytic continuation of the disk-to-polygon map across each side of the polygon (p. 108), the same continuation this volume's section 22.8 invoked at second hand from Driscoll and Trefethen.*

*Two further passages bear directly on this chapter's own claims. First (pp. 109–110), Schwarz derives the local exponent law at a corner in general form: a wedge of angle $\alpha\pi$ at the origin, mapped conformally to a flat half-plane boundary, is realized near the corner by $\varrho = u^{1/\alpha}$ composed with an otherwise-ordinary analytic map, so that the target coordinate behaves as $t - t_0 \sim C(u - u_0)^\alpha$ near a corner of interior angle $\alpha\pi$ — exactly the correspondence section 22.9's "from the target polygon's interior angle" argument used, there cited from Driscoll–Trefethen's modern restatement, to fix $\beta_k = 1/2$ at each of Peirce's four singular points. Schwarz's own p. 110 — independently flagged by Ringleb (1927) as the single page his own paper on multivaluedness depends on, per the earlier todobox above, now confirmed rather than only guessed at — is where Schwarz observes that the quantity $\frac{d}{dt}\log\frac{du}{dt}$ is invariant under replacing $u$ by any Möbius transform $C_1u + C_2$: the combination is unaffected by the target figure's arbitrary placement and scale, which is exactly why it, rather than $u$ or $du/dt$ themselves, is the right object to expand near a singular point. A few pages later (p. 116) Schwarz forms the second-order combination $\frac{d^2}{dt^2}\log\frac{du}{dt} - \frac{1}{2}\left(\frac{d}{dt}\log\frac{du}{dt}\right)^2$, calling it $\Psi(u,t)$ and showing it too is invariant under replacing $u$ by $C_1u + C_2$ (and, taken with $u, t$ interchanged, satisfies a change-of-variable law under replacing $t$ itself by a rational function) — this is, in Schwarz's own original notation and construction, the object now called the Schwarz derivative or Schwarzian derivative $S(u,t)$, introduced here for exactly the purpose (characterizing circular-arc-polygon maps via a linear ODE with rational coefficients, reducing to the hypergeometric equation for the triangle case) that the name still commemorates today.*

*Second, the paper's own worked instance of the square-to-disk problem (pp. 110–112) is the specific case this volume's section 22.3 draws its elliptic-function formula from at second hand. With singular points at $t = -1, 0, 1$ (equivalently $t = \infty$) and $\alpha = \frac{1}{2}$ at each, Schwarz's invariant reduces to the rational function $-\frac{1}{2}\left(\frac{1}{t+1} + \frac{1}{t} + \frac{1}{t-1}\right)$, and a single integration gives, in Schwarz's own notation,*

$$u \;=\; C_1 \int_0^t \frac{dt}{\sqrt{4t(1-t^2)}},$$

*identified explicitly (p. 111) as a lemniscatic integral; a further substitution $s = (t-i)/(t+i)$ carries the half-plane to the unit disk, and the pair $u = \int_0^s ds/\sqrt{1-s^4}$, $s = \operatorname{sinam}(u)$ with $k = \sqrt{-1}$ (p. 112) maps the unit disk conformally onto a square with vertices $K, Ki, -K, -Ki$, where $K = \int_0^1 ds/\sqrt{1-s^4}$ is the lemniscate constant — Schwarz includes a numerical table of* sinam *values on the tenths-of-K grid and a drawing of the two grids (square and disk, superimposed at 45°) that this volume's own reflection-principle argument in section 22.8 echoes structurally, though built there from a different elliptic function (Peirce's own, via the quincuncial substitution) rather than reproduced from Schwarz's page. This confirms, from the primary source rather than a modern secondary account, that the disk-to-square construction section 22.3 cites as "a specific instance of the general Schwarz–Christoffel construction" really is exactly Schwarz's own worked example, down to the choice of singular points and the lemniscatic (as opposed to a general elliptic) modulus forced by the square's fourfold symmetry — and, separately, that section 22.11's finding (Peirce cites no source for his own instance of this same*

**Epistemic Notice**

This reading closes the access half of section 22.12's item for Schwarz specifically; Christoffel's 1867 paper (section 22.12) remains reached but unread, still behind the same scanned-Italian-typesetting barrier met for Krüger's monograph, and is not affected by this pass. Nothing this section reads changes any result already proved in this chapter: propositions 22.1, 22.2 and 22.6's arguments were already correct on Driscoll and Trefethen's modern restatement of the same facts, and the primary source confirms rather than corrects them. What changes is provenance — the $\alpha = 1/2$ exponent law and the lemniscatic-integral formula this chapter draws on are now traced to Schwarz's own pp. 109–112, not merely to a textbook's account of them — and one small addition of intellectual-historical interest not previously recorded in this volume: Schwarz's own $\Psi(u, t)$ (p. 116) is, in essence and construction though not in name, the Schwarzian derivative, introduced in this very paper for the circular-arc-polygon problem, a fact worth noting since this volume has cited the Schwarz–Christoffel construction repeatedly without previously registering that the paper establishing it also originates a second, independently famous piece of complex-analytic machinery.

**Key Result**

*Three of table 26.1's originally flagged gaps are now closed (aphylactic family by chapter 19, interruption itself by chapter 20, aphylactic-and-interrupted by chapter 21); this chapter examines, and leaves open with a sharper diagnosis, the fourth (conformal-and-interrupted). Table 26.1 gains an eighth row from this chapter regardless of that outcome, since a conformal, contiguous case built from a polyhedral skeleton is itself a new data point about the deployment and domain axes even where it does not move the continuity axis, and the cone-point identification of section 22.8, sharpened in section 22.9 from a qualitative type to a computed exponent ($\alpha = 1/2$ at each of the four points, confirmed by two independent arguments and by the resulting $4\pi$ global defect check) is new to this volume's account of the family axis's boundary behavior at singular points generally.*

## 22.13 The 1877 Coast Survey report, located: the appendix is two pages, not the volume as a whole

Section 22.11's todobox left the 1877 United States Coast Survey annual report — the presumed source of section 22.2's 9 %/13 %/50 % comparison, distinct from the 1879 journal paper already read in full — narrowed to a named, reachable NOAA digitization project but not yet to a specific retrievable file. A further pass located the specific scan directly: NOAA's Central Library rescue-imaging project hosts the volume as `CSC-0026.PDF`, and its own front matter, table of contents, and alphabetical index are legible in ordinary extracted text rather than only as page images, which is itself a better access outcome than section 22.12's scanned-Italian and OCR-limited German cases.

**Key Result**

*[The appendix's own pagination, read from the volume's own front matter rather than from a secondary bibliography] The report's table of contents lists its fifteen appendices by title and page range, and its own alphabetical index cross-references the item twice, under both "Peirce, C. S., Assistant" and "Quincuncial projection of the sphere": in both places, and in the contents listing itself, Appendix No. 15, "A Quincuncial Projection of the Sphere. By Chas. S. Peirce, Assistant," is given the page range **191, 192** — two pages, not three or four. This settles, from the primary volume's own front matter rather than from a secondary citation of it, a discrepancy this chapter's own bibliography had not previously flagged: a modern GIS-vendor documentation page, consulted only for the 1879 journal citation and not otherwise relied on by this chapter, separately gives the 1877 appendix's range as pp. 191–194, four pages rather than two. That four-page figure is not what the 1877/1880 report's own contents and index record. A separately paginated plate (the map itself, unnumbered within the body pagination, listed elsewhere in the same report as one of a run of numbered progress sketches) accompanies the two-page text, which is consistent with, and independently corroborates, section 22.11's own guess — made there about the different, 394-versus-397 discrepancy in the 1879 journal printing — that an accompanying map plate is the recurring source of these one-or-two-page mismatches across this chapter's bibliography, now confirmed in a second, independent instance rather than only conjectured for the first.*

**Todo — verify primary source**

Section 22.13 closes the specific-document half of section 22.11's todobox (the 1877 report is now identified down to its exact host file and exact pagination) but not the substantive half: this pass's retrieval of `CSC-0026.PDF` returned the volume's front matter, letter of transmittal, abstract of contents, and full alphabetical index in legible extracted text, but was exhausted before reaching the body pagination in the 190s where Appendix No. 15 itself sits, so the two-page appendix text — and with it, whether section 22.2's 9 %/13 %/50 % comparison actually appears there rather than in some other, still-unidentified source — remains unread. What has changed in degree is substantial: the item is narrowed from "an unidentified report, in an unidentified collection" to "a named two-page appendix in a named, already-open PDF, at a known point in its pagination not yet reached by this pass's retrieval," the same kind of narrowing-without-closing section 22.12 recorded for Schwarz's paper before a targeted page-range fetch finally closed it.

**Todo — verify primary source**

A further pass in this same pass (v78) tried the natural next step on section 22.13's own terms: rather than re-fetching `CSC-0026.PDF` at a larger length allowance (the approach already shown, for a different file, not to help by section 22.12's repeated Schwarz–Christoffel attempts), it tried reaching the same appendix through NOAA's own bibliographic index of Coast Survey appendices, on the reasoning that a per-appendix index entry might link directly to a page-addressable view rather than the flat front-to-back PDF already narrowed to a known-but-unread pagination. That index (`library.oarcloud.noaa.gov/docs.lib/htdocs/rescue/cgs/CGS_appencies/chronological.html`) is a single continuous listing spanning 1844–1910, long enough that the specific 1877 entries fell outside what one retrieval reached, and a further, more targeted search for the exact NOAA-hosted 1877 filename did not surface it independently either. This closes nothing section 22.13's own todobox left open — the two-page appendix itself remains unread on exactly the same terms — but records a second access route now tried and not yet productive, alongside the direct large-length-allowance refetch already known not to help, so a future pass has two ruled-out approaches to avoid repeating rather than one.

**Todo — verify primary source**

A still later pass (5 Aug 2026) tried a third direction, external to NOAA's own site entirely: rather than fetching `CSC-0026.PDF` or its index again, it searched for existing citations of the file itself, on the reasoning that a scholar who had already opened the PDF successfully might record, in citing it, exactly the offset this chapter still lacked. This succeeded in a specific and useful way. Peirce's own dedicated secondary bibliography (the Arisbe project, hosted at `cspeirce.com`) cites this identical file with a page-anchored URL, `docs.lib.noaa.gov/rescue/cgs/001_pdf/CSC-0026.PDF#page=215`, and separately confirms the 1879 journal article's own pagination (pp. 394–397) already on record in this chapter's bibliography, consistent with section 22.13's own finding and adding nothing new on that point. The page-215 figure, however, is new and is a different kind of narrowing than section 22.13's: that result gave the appendix's pagination within the 1877 report's own printed body (pp. 191–192), which is not the same number as the offset within the scanned PDF file that a targeted fetch actually requires, since the scan's own front matter, letter of transmittal, and index shift the two out of alignment by several dozen pages; this result supplies exactly the PDF-internal offset, sourced from a third party's prior successful access rather than computed from this chapter's own reading of the front matter. A direct attempt to fetch that exact anchored URL was refused outright by `docs.lib.noaa.gov` (a robots-disallowed response) — a different failure mode from anything previously recorded in this section: not a length-insensitive truncation of the kind section 22.12 diagnosed for the Internet Archive Crelle volume, but a host that declines automated retrieval regardless of what is requested or how, which rules out any further attempt at that specific host and domain rather than narrowing what request might work there. The outstanding step for a future pass is accordingly the same targeted retrieval redirected to the `library.oarcloud.noaa.gov` mirror already used successfully for this file's front matter in section 22.13, now aimed at its internal page 215 rather than at the file's start, which neither this pass nor either of the two preceding it had a page number to aim at.

### 22.13.1 A further pass (6 Aug 2026): the oarcloud mirror reached, and a new, more precise diagnosis than a robots block

A further pass took up exactly the outstanding step the previous todobox left: a targeted retrieval against `library.oarcloud.noaa.gov` rather than the robots-blocked `docs.lib.noaa.gov` host. Before attempting the fetch itself, this pass first cross-checked section 22.13's file identification independently, by reconstructing the fond's own year-to-file mapping from the rescue-imaging project's sequential numbering (1852 = `CSC-0001`, 1853 = `CSC-0002`, and so on): this confirms, from the collection's own internal structure rather than from Arisbe's citation alone, that 1877 corresponds to `CSC-0026` exactly as section 22.13 already had it.

**Key Result**

*[The oarcloud mirror, fetched: a fourth host tried for this item, and the first not refused outright]* *The mirror, at `library.oarcloud.noaa.gov/docs.lib/htdocs/rescue/cgs/001_pdf/CSC-0026.PDF`, is not disallowed by robots policy, unlike the `docs.lib.noaa.gov` host the previous pass's fetch was refused by. The fetch succeeded and confirms, from this fourth independent route to the same file, exactly what section 22.13 already had from the volume's own front matter: an index entry for "Appendix No. 15 … pp. 191, 192 … A QUINCUNCIAL projection of the sphere … By Chas. S. Peirce," and a second, alphabetical-index cross-reference under Peirce's own name. Nothing about the appendix's identification or pagination changes; what changes is that a working, non-blocked route to the file now exists on record, where the three routes tried before it (the direct `docs.lib.noaa.gov` front-to-back fetch, the same host's page-anchored URL, and the govinfo.gov Serial Set package) were each refused, exhausted, or unreachable by script for a different reason apiece.*

**Todo — verify primary source**

This reachable route does not, however, close the item, and the way it fails to is new rather than a repetition of any earlier diagnosis in this section. The extracted text jumps directly from the front matter and table of contents to the volume's closing alphabetical index, skipping the entire intervening body and appendices — pp. 1–192, exactly the range containing Appendix No. 15 itself — rather than truncating partway through it, the failure mode section 22.12 diagnosed three times over for the Crelle volume. A repeat fetch with the token limit raised to 200,000 (an order of magnitude past this pass's default) returned a result identical byte for byte to the first attempt, ruling out request length as the cause in the same manner section 22.12's own repeated attempts eventually ruled it out for Schwarz, but by a single confirming fetch here rather than three. The diagnosis this leaves is structural rather than access-related: the front matter and the index carry a machine-readable OCR text layer, but the body and appendices apparently do not, consistent with that portion of the scan being image-only pages rather than OCR-processed ones. No further request to this host, at any length, is expected to change this; the only route left untried is optical character recognition run directly against the page images themselves, a step outside what this pass's tools can perform against a host external to this volume's own working environment. The item is accordingly narrower and more precisely stated than at any point above — a named two-page appendix, in a named file, on a reachable host, at a diagnosed and specific point of failure — but remains open on those terms, alongside appendix A.8.3's Maurer item in appendix A as one of this volume's two remaining primary-source gaps with a clear, if not yet successful, next step.

# Chapter 23

# This Volume's Own Answer: An Interrupted Construction, Conformal on Every Gore

## 23.1 Returning to the question chapter 22 left open

Section 22.10's closing disclaimerbox left a specific question open: whether a genuinely interrupted conformal projection can be constructed at all, given that Peirce's own route to conformality (a single global conformal map via Riemann's mapping theorem) produces a seamless, contiguous domain by its very construction, in apparent tension with the route chapters 20 and 21 used to achieve interruption (cutting a polyhedral net along chosen edges and arranging the pieces with gaps). This chapter answers the question, using a construction assembled entirely from results this volume has already proved rather than any new cited material: yes, in the specific sense made precise below, and the construction that shows it needs nothing beyond proposition 22.2 (proved in chapter 22) and the icosahedral net (used already in chapter 21).

## 23.2 The construction: stereographic projection, applied gore by gore

**Definition 23.1** (Interrupted stereographic-icosahedral construction)**.** Let $\mathcal{I}$ be the regular icosahedron of definition 21.1, with its twenty spherical triangular faces. For each face $F$, let $S_F$ denote the point on the unit sphere antipodal to $F$'s centroid direction $\widehat{n}_F$, and let $\pi_F$ be the stereographic projection of the sphere from $S_F$ onto the plane tangent at $\widehat{n}_F$ (the construction of proposition 22.2, recentred at $F$). The *interrupted stereographic-icosahedral projection* sends a point $P$ in $F$'s spherical triangle to $\pi_F(P)$, and arranges the twenty resulting plane triangles according to $\mathcal{I}$'s own combinatorial net, cut (rather than folded) along the same edges section 21.7 left as an open combinatorial choice for the historical Fuller map.

This is, deliberately, the same combinatorial skeleton definition 21.1 used for the Fuller-proxy construction — twenty faces, the same net, the same choice of which edges to cut — with exactly one change: definition 21.1's gnomonic projection from the sphere's *center*, replaced here by a stereographic projection from each face's own *antipodal* point. Proposition 21.3 showed the first choice is aphylactic away from each face's own centroid; the next proposition shows the second choice is not.

**Proposition 23.2** (Every gore of definition 23.1 is conformal, in full)**.** *For each face $F$, the restriction of $\pi_F$ to $F$'s spherical triangle is conformal at every point of that triangle, not merely at its centroid.*

*Proof.* Proposition 22.2 proved $h(c) = k(c) = \sec^2(c/2)$ for stereographic projection at *every* angular distance $c \in [0, \pi)$ from the projection's own tangent point, with no restriction to a neighborhood of $c = 0$: the identity $h(c) = k(c)$ was shown to hold identically, for the full range of $c$ over which the projection is defined. A face $F$'s spherical triangle, however large, occupies angular distances from $\widehat{n}_F$ strictly less than $\pi$ (indeed, by section 21.4's numerical computation, at most the icosahedron's own vertex angular radius, approximately 37.38°, far short of the 180° at which stereographic projection's domain ends), so $F$'s triangle lies entirely within $\pi_F$'s domain of conformality, and proposition 22.2's identity $h(c) = k(c)$ applies at every one of its points without exception. □

**Key Result**

*Proposition 23.2 is a sharp contrast with proposition 21.3's finding for the same combinatorial construction built from gnomonic rather than stereographic projection: where the Fuller-proxy map of chapter 21 is aphylactic at every point except twenty isolated centroids, this chapter's stereographic-icosahedral map is conformal at* every *point of every one of its twenty gores, with no exception at all within a face's interior. The reason is exactly section 22.10's keyresult already noted in passing: stereographic projection's scale factors are equal identically in c, not merely at $c = 0$ the way a general azimuthal projection's scale factors might happen to agree only at the origin, and this global (rather than merely local) conformality is what survives the restriction to a single face unchanged.*

## 23.3 Interruption, exactly as before

Definition 23.1's domain is disconnected into twenty separate planar gores by exactly the same combinatorial cutting section 21.7 used for the Fuller-proxy construction — the same net, the same chosen edges. Nothing about replacing gnomonic with stereographic projection changes this: interruption is a fact about which edges are cut and which are folded, entirely independent of which azimuthal projection is used within each face.

**Key Result**

*By definition 26.1's continuity axis, definition 23.1's construction is interrupted in precisely the sense chapters 20 and 21 established it: a domain cut by design into separate gores, not a single connected region. Combined with proposition 23.2, this construction is simultaneously conformal (on every point of every gore, not merely at isolated centroids) and interrupted — the family-by-continuity combination sections 21.9 and 22.10 both flagged as unfilled.*

## 23.4 What this closes, and what it does not

**Epistemic Notice**

This construction answers section 22.10's open question in a specific, narrow sense that this section states carefully rather than overclaims. It is not a historical case study: unlike every other entry in table 26.1, definition 23.1 was not devised by a named historical cartographer and has no publication history for this volume to research: it is assembled, in this chapter, entirely from two pieces already proved elsewhere in this volume (proposition 22.2 and the icosahedral net of definition 21.1), and its only claim to novelty is the specific combination, not either ingredient. Section 22.10's suggestive tension — that Peirce's particular route to conformality (one global Riemann map) resists interruption by its own construction — is not thereby refuted; it is sidestepped, by using a different, more elementary route to conformality (per-gore stereographic projection, not a single global elliptic-function map) that has no reason to produce a single connected image in the first place. Whether every *historically named* conformal projection resists interruption for a structural reason, or whether this is an artifact of which historical examples this volume happened to examine, remains exactly as open as section 22.10 left it; what this chapter shows is only that the family-by-continuity cell itself is not empty, answered by a construction rather than by an argument that no obstruction exists.

**Key Result**

*Table 26.1 can now record all four family-by-continuity combinations as populated: aphylactic/-contiguous (Robinson), equal-area/interrupted (Goode), aphylactic/interrupted (Fuller), and now conformal/interrupted, by the construction of this chapter rather than by a named historical map. Chapter 22's own conformal/contiguous case (alongside Mercator and Gauss–Krüger) rounds out the sixth of the six family-by-continuity cells, leaving none of the three-family-by-two-continuity grid unpopulated, though — as throughout this volume's use of the word — "populated" records a witnessing construction for a cell, not an exhaustive census of every projection that could occupy it, and the eighteen cells of the fuller twenty-four-cell grid (crossing deployment and domain as well) remain almost entirely open, exactly as table 26.1's own caption has said throughout.*

# Chapter 24

# Three More Cases: A National Geographic Lineage, and a Controversy About What "Best" Means

## 24.1 Why these three, together

This chapter was added after chapters 19 to 23 and after chapter 26's own synthesis, in the same spirit section 26.5 already used to describe chapters 8 to 10, 15 and 17: as an extension that deepens Part III's historical record without reopening chapter 26's own count-keeping, which this chapter leaves untouched exactly as section 26.5 left it. Section 24.5 below is explicit about why, and about the one place this chapter's own findings do bear on that count.

Three named constructions are treated together here because they are connected by two independent threads, not one. The first is institutional and chronological: Alphons van der Grinten's projection (section 24.2) was the National Geographic Society's standard world-map projection from 1922 until 1988, when it was replaced by Robinson's (chapter 19's own subject), which in turn was replaced in 1998 by Oswald Winkel's Winkel Tripel (section 24.3) — three successive aphylactic choices by the same institution across seventy-six years, each adopted specifically because its predecessor's compromise was judged, by the cartographic standards of its day, no longer the best available one. The second thread is a genuine controversy, not merely a series of institutional replacements: the projection popularly known as Gall–Peters (section 24.4) became, in the 1970s and 1980s, the subject of a public dispute about whether area-accurate maps are more honest than shape-accurate ones, a dispute this chapter reports on section 24.4.2's terms — presenting the documented positions on each side, not adjudicating which is correct, consistent with section 1.4's own commitment to fair presentation of contested questions this series does not treat as settled by mathematics alone.

## 24.2 Van der Grinten (1898/1904): the circle that came before Robinson

Alphons J. van der Grinten devised his best-known projection in 1898 and patented it, together with three related constructions, in 1904 (US Patent #751,226).[1] The National Geographic Society adopted it for its reference world maps in 1922, a choice that raised the projection's visibility considerably beyond its original patent filing, and used it until 1988, when it was replaced by chapter 19's Robinson projection — the first link in this chapter's institutional chain.

Unlike every projection chapters 14 and 16's case studies treated, van der Grinten's construction was not derived from a stated invariant (conformality, equal area, equidistance along some locus) at all: van der Grinten specified it geometrically, as a compass-and-straightedge construction projecting the whole sphere into a circle, with the equator and central meridian as straight diameters and every other meridian and parallel drawn as a circular arc.[2] The algebraic formulas usually cited today were published only decades later; this chapter uses the version given by Snyder's 1987 USGS working manual,[3] cross-checked identically across two independent secondary transcriptions before being trusted:

$$\begin{aligned} x &= \frac{\pm\pi\left[A(G-P^2)+\sqrt{A^2(G-P^2)^2-(P^2+A^2)(G^2-P^2)}\right]}{P^2+A^2}, \\ y &= \frac{\pm\pi\left[PQ-A\sqrt{(A^2+1)(P^2+A^2)-Q^2}\right]}{P^2+A^2}, \end{aligned} \tag{24.1}$$

where $x$ takes the sign of $\lambda-\lambda_0$, $y$ takes the sign of $\varphi$, and

$$A=\frac{1}{2}\left|\frac{\pi}{\lambda-\lambda_0}-\frac{\lambda-\lambda_0}{\pi}\right|, \quad \theta=\arcsin\left|\frac{2\varphi}{\pi}\right|, \quad G=\frac{\cos\theta}{\sin\theta+\cos\theta-1},$$
$$P=G\left(\frac{2}{\sin\theta}-1\right), \quad Q=A^2+G,$$

with the degenerate cases $\varphi=0$ (giving $y=0$, $x=R(\lambda-\lambda_0)$, the equator scale-true by construction) and $\lambda=\lambda_0$ or $\varphi=\pm\pi/2$ (giving $x=0$) handled separately, exactly as eq. (24.1)'s source states.

**Proposition 24.1** (Van der Grinten is aphylactic, confirmed by direct computation)**.** *This chapter computed the local scale factors $(h,k)$ and the areal scale factor $hk\sin\omega'$ (the Jacobian-determinant form chapter 3's own framework already uses, evaluated by a central-difference numerical gradient of eq.* (24.1) *rather than an analytic derivative, consistent with chapter 19's own numerical treatment of a construction with no simple closed-form derivative)*

[1]The 1898 devising date, the 1904 patent and its number, and the four-projection patent family are corroborated across multiple independent secondary and tertiary sources consulted 5 Aug 2026 (including a USGS-sourced secondary account and several independent encyclopedic entries), all citing Snyder's 1993 *Flattening the Earth*, pp. 258–262, and Snyder & Steward's 1989 USGS Bulletin 1856 bibliography, p. 94, as the underlying primary-adjacent record; neither has been consulted directly in this pass, and the todobox below records this.

[2]The geometric (compass-and-straightedge) origin of the construction, contrasted explicitly with the algebraic formulas published decades later, is reported consistently across the secondary literature, including an account crediting the algebraic formulation to O'Keefe & Greenberg (1977) and to Snyder (1979b), neither consulted directly here.

[3]J. P. Snyder, *Map Projections — A Working Manual*, USGS Professional Paper 1395, 1987, pp. 239–242; cross-checked against two independent secondary transcriptions (a Wolfram MathWorld entry and an independent formula-archival site) before use here, following exactly the cross-checking discipline chapter 19's own numerical table already applied. Not consulted directly in this pass; see the closing todobox.

*at four points spanning the mapped hemisphere. Both the shape factor* $\max(h,k)/\min(h,k)$ *and the areal factor* $hk\sin\omega'$ *(evaluated as the ratio* $\sqrt{EG-F^2}/\cos\varphi$ *of the plane metric to the sphere's own area element) depart from* 1 *at every point tested away from the origin, growing sharply toward the mapped hemisphere's edge:*

| $(\lambda,\varphi)$ | *shape factor* | *area factor* |
|---|---|---|
| $(30°, 0°)$ | 1.000 | 1.000 |
| $(60°, 20°)$ | 1.012 | 1.121 |
| $(90°, 40°)$ | 1.052 | 1.637 |
| $(120°, 60°)$ | 1.137 | 3.492 |

*confirming, by direct computation rather than by the informal "it looks like it distorts everything a bit" characterization the secondary literature gives, that van der Grinten's construction is genuinely aphylactic in definition 4.4's sense — neither conformal nor equal-area — and that the extreme polar-region distortion the secondary literature reports*[4] *is a real, computed feature of eq.* (24.1)*, not only a qualitative visual impression: the area factor alone grows by a factor of nearly* 3.5 *between* $(30°, 0°)$ *and* $(120°, 60°)$*, the steepest growth this volume's Part III has computed for any single construction over a comparable angular range.*

## 24.3 Winkel Tripel (1921): an average of two other projections

Oswald Winkel proposed three projections in a single 1921 paper, each constructed as the arithmetic mean, coordinate by coordinate, of the equirectangular projection and one other named projection: the sinusoidal for Winkel I, the Mollweide for Winkel II, and the Aitoff projection for Winkel III — the construction now universally known simply as the Winkel Tripel.[5] "Tripel" is German for "triple," naming Winkel's goal of moderating three kinds of distortion — area, direction, and distance — at once, rather than eliminating any one of them exactly.

The construction is, unusually for a named projection this series has treated, defined entirely in terms of two others rather than by an independent formula. Writing $(x_E, y_E)$ for the equirectangular projection with standard parallel $\varphi_1$ (the same construction section 24.4 below uses for a different $\varphi_1$),

$$x_E = \lambda\cos\varphi_1, \qquad y_E = \varphi,$$

and $(x_A, y_A)$ for the Aitoff projection,

$$\alpha = \arccos(\cos\varphi\cos(\lambda/2)), \qquad x_A = \frac{2\cos\varphi\sin(\lambda/2)}{\operatorname{sinc}\alpha},$$
$$y_A = \frac{\sin\varphi}{\operatorname{sinc}\alpha}, \qquad \operatorname{sinc}\alpha := \frac{\sin\alpha}{\alpha}\ (\operatorname{sinc}0 := 1),$$

[4]The specific characterization "extreme distortion in polar regions," and the consequent practice of published van der Grinten maps rarely extending past roughly the northern shores of Greenland, is reported across the secondary literature, including a USGS-sourced secondary account of van der Grinten's own working method.

[5]The three-projection 1921 paper and the specific pairing (equirectangular with sinusoidal, Mollweide, and Aitoff respectively) is corroborated by an independent secondary encyclopedic account of the Winkel projection family, cross-checked against Winkel's own promotional site (winkel.org, itself a secondary rather than primary source for the 1921 paper) and against Snyder's 1993 *Flattening the Earth*, pp. 231–232, as cited by the secondary literature; neither Winkel's 1921 paper nor Snyder's own text has been consulted directly in this pass.

the Winkel Tripel coordinates are the arithmetic mean $x = (x_E + x_A)/2$, $y = (y_E + y_A)/2$.[6] National Geographic adopted the Winkel Tripel in 1998, with $\varphi_1 = 50°28'$, replacing Robinson — the second and, to date, final link in this chapter's institutional chain.

**Proposition 24.2** (The Winkel Tripel, confirmed aphylactic by direct computation of its averaged construction)**.** *This chapter implemented the averaging construction directly and computed local scale factors by the same numerical-gradient method proposition 24.1 used:*

| $(\lambda, \varphi)$ | $h$ | $k$ | *area factor* $hk$ |
|---|---|---|---|
| $(0°, 0°)$ | 1.000 | 0.818 | 0.818 |
| $(30°, 0°)$ | 1.006 | 0.818 | 0.823 |
| $(0°, 30°)$ | 1.000 | 0.891 | 0.891 |
| $(60°, 45°)$ | 1.055 | 0.992 | 1.013 |
| $(90°, 60°)$ | 1.187 | 1.175 | 1.219 |

*Neither $h = k$ (conformality) nor $hk = 1$ (equal area) holds at any of the five points tested, confirming by direct computation — as opposed to the averaging construction's own suggestive name, which does not by itself entail either property fails everywhere — that the Winkel Tripel is genuinely aphylactic. This is worth stating explicitly rather than assuming: an arithmetic mean of two projections' coordinates has no a-priori reason to inherit a scale-factor identity from either component (the mean of two functions satisfying $h = k$ need not itself satisfy $h = k$, since the scale factors are computed from derivatives of the mean, not the mean of the derivatives' own ratios), and proposition 24.2's table confirms this generic expectation rather than any special structure of the Winkel construction specifically.*

**Epistemic Notice**

This chapter's Winkel Tripel numbers are a first-principles computation from eq. (24.2)'s reconstructed formula, not a reproduction of any published distortion table; they have not been cross-checked against Goldberg & Gott's (2007) published *Cartographica* study of the projection's flexion and skewness properties, cited by chapter 24's own secondary sources as the standard modern reference for exactly this kind of computation but not consulted directly in this pass. The qualitative conclusion — aphylactic, with low but nonzero distortion of every kind, matching the "Tripel" name's own stated goal — is consistent with every secondary source consulted, several of which report, without this chapter re-deriving the specific figure, that the 2007 Goldberg–Gott study found the Winkel Tripel's overall distortion favorably comparable to, and by some measures better balanced than, chapter 19's own Robinson projection — the two aphylactic projections that were, in fact, the direct institutional successor and predecessor at the National Geographic Society.

$$x = \frac{1}{2}\left(\lambda \cos\varphi_1 + \frac{2\cos\varphi \sin(\lambda/2)}{\operatorname{sinc}\alpha}\right), \qquad y = \frac{1}{2}\left(\varphi + \frac{\sin\varphi}{\operatorname{sinc}\alpha}\right), \qquad \varphi_1 = 50°28'. \tag{24.2}$$

[6] Both component formulas (equirectangular with adjustable standard parallel, and Aitoff via the halved-longitude azimuthal-equidistant construction) are standard, textbook-level closed forms independently corroborated across multiple sources consulted 5 Aug 2026, including a physics preprint's appendix and an independent secondary cartographic encyclopedia entry for each of the two component projections; the averaging construction itself, and Winkel's own choice of $\varphi_1 = 50°28'$ for the equirectangular component, are reported identically by Esri's ArcGIS technical documentation and by an independent cartographic-history account, and this chapter reconstructs eq. (24.2) below from these two verified components rather than transcribing a rendered formula directly from Wikipedia's own article, whose formula images this chapter's fetch tooling could not extract as text; the underlying facts asserted (the averaging construction, the two named components, and Winkel's $50°28'$ choice) are independently corroborated by the prose of every source consulted, only the rendered equation itself was not extractable.

## 24.4 Gall (1855) and Peters (1967/1973): one projection, two names, and a genuine equal-area proof

The projection popularly called Gall–Peters is, mathematically, the simplest of this chapter's three cases: a cylindrical equal-area projection, in the specific one-parameter family already implicit in chapter 4's own trade-off framework, with its single free parameter (the standard parallel $\varphi_1$, the latitude at which the projection is locally undistorted in shape) fixed at $\varphi_1 = 45°$:

$$x = \lambda \cos \varphi_1, \qquad y = \frac{\sin \varphi}{\cos \varphi_1}, \qquad \varphi_1 = 45°. \tag{24.3}$$

**Proposition 24.3** (Exact equal-area property, proved in closed form)**.** *Equation* (24.3) *is equal-area at every point, exactly, not merely to within numerical tolerance. Since* $\partial x/\partial \varphi = 0$ *and* $\partial y/\partial \lambda = 0$ *identically, the plane metric's cross term vanishes and the areal scale factor is simply the product of the two partial derivatives' magnitudes, divided by the sphere's own area element* $\cos \varphi$*:*

$$\textit{area factor} = \frac{1}{\cos \varphi} \left| \frac{\partial x}{\partial \lambda} \right| \left| \frac{\partial y}{\partial \varphi} \right| = \frac{1}{\cos \varphi} \cdot \cos \varphi_1 \cdot \frac{\cos \varphi}{\cos \varphi_1} = 1,$$

*identically in both* $\lambda$ *and* $\varphi$*, for every choice of* $\varphi_1$ *— not only* $45°$*. This is the same one-line cancellation chapter 14's own treatment of Lambert's azimuthal projection already exhibited for a different family member; table 26.1's existing equal-area/global/continuous/contiguous row (Lambert, 1772) and this section's Gall–Peters entry are, in this precise sense, two different closed-form witnesses to the identical algebraic mechanism, a cylindrical member and an azimuthal member of the same provably-exhaustive equal-area family corollary 4.5 names.*

This chapter's own numerical check, at $\varphi = 0°, 30°, 45°, 60°, 80°$, reproduces proposition 24.3's exact identity to $3 \times 10^{-7}$ relative error (limited only by the central-difference step size used, not by any property of eq. (24.3) itself), and confirms the standard-parallel property directly: $h = k = 1$ exactly at $\varphi = 45°$ (no distortion of any kind there), with $h$ falling and $k$ rising away from it in exactly the reciprocal pattern equal-area cylindrical projections must show, reaching $h = 0.246$, $k = 4.072$ at $\varphi = 80°$ — the same severe high-latitude shape distortion (elongation along parallels, compression along meridians) that section 24.4.2 below reports as the central cartographic objection to the projection's use for general-purpose world maps.

### 24.4.1 Two names, and how the second one was affixed

James Gall, a Scottish clergyman and amateur cartographer, presented the construction — calling it his "orthographic" projection, despite bearing no relation to the true orthographic projection — at the 1855 Glasgow meeting of the British Association for the Advancement of Science, and published it formally thirty years later, in 1885, in the *Scottish Geographical Magazine*.[7] The construction attracted little attention for roughly a century. In 1967, the German historian and filmmaker Arno Peters devised, independently by his own account, a projection with the same mathematical properties, and presented it in 1973 at a press

[7]The 1855 presentation, the specific venue (Glasgow meeting of the British Association for the Advancement of Science), and the 1885 *Scottish Geographical Magazine* publication are corroborated identically across multiple independent secondary and tertiary sources consulted 5 Aug 2026; none of Gall's own 1855 or 1885 texts has been consulted directly in this pass, and the closing todobox records this.

conference in Bonn as a new invention, without reference to Gall's prior work.[8] Peters's own published description of his projection's construction, taken literally, implies standard parallels at 46°02′ rather than 45°; the secondary literature reports that Peters's accompanying text makes clear he intended 45°, so that his construction and Gall's are, as published, identical (the discrepancy negligible at world-map scale in any case).[9] The composite name "Gall–Peters" was coined, the secondary literature reports, by Arthur H. Robinson — chapter 19's own subject — in a 1986 American Cartographic Association pamphlet, a detail worth noting for the same reason section 24.1 opened this chapter with an institutional lineage: Robinson is, on this record, not only the direct institutional predecessor of Winkel Tripel and successor of van der Grinten at the National Geographic Society, but also the person who fixed, in print, the name under which this section's third case is now most widely known.[10]

### 24.4.2 The controversy, presented on its own terms

Peters's public case for his projection, as reported consistently across the secondary literature, rested on a specific empirical observation about the Mercator projection's areal distortion (already established in this volume's own terms by section 14.2's keyresult) combined with a claim about its social and political consequences: because much of the world's less industrialized territory lies at low latitudes, where Mercator's areal exaggeration is smallest, and much of the industrialized world lies at higher latitudes, where that exaggeration is largest, Mercator world maps — Peters argued — systematically understate the visual prominence of the developing world relative to the industrialized one, and an equal-area alternative corrects this specific bias.[11] This argument found real institutional traction: maps based on the projection have been promoted by UNESCO and used in some British schools, and the projection reached a wide popular audience, including a fictionalized dramatization of essentially this argument in a widely-seen American television drama's 2001 episode.[12]

The professional cartographic community's response, as reported across the secondary literature, was not a rejection of equal-area mapping as such — proposition 24.3's own proof shows the underlying mathematics is exactly as sound as Lambert's 1772 azimuthal equal-area projection, already a settled case in this volume by chapter 14 — but a specific two-part objection: first, that equal-area cylindrical projections were neither new nor uniquely due

[8]The independence claim, the 1967 devising date, the 1973 Bonn press conference, and the absence of any acknowledgment of Gall's prior work in Peters's own presentation are reported consistently across multiple independent secondary sources consulted 5 Aug 2026, several of which frame Peters's lack of formal cartographic training and consequent unfamiliarity with the existing literature as the likeliest explanation, rather than deliberate misattribution; this chapter reports the claim and the surrounding record without adjudicating Peters's own state of knowledge, which no source consulted here settles beyond secondary characterization.

[9]The 46°02′ literal-versus-45°-intended discrepancy in Peters's own published description is reported identically across multiple independent secondary sources; not independently verified against Peters's own text in this pass.

[10]The 1986 American Cartographic Association pamphlet attribution to Robinson is corroborated across multiple independent secondary sources, including direct citation of the pamphlet's existence (*Which Map Is Best?*, American Cartographic Association Special Publication No. 1, 1986); the pamphlet itself has not been consulted directly in this pass.

[11]This characterization of Peters's argument is reported consistently across the secondary literature consulted 5 Aug 2026, including the specific comparison, repeated across multiple independent sources, of Greenland's appearance under Mercator (visually comparable to or larger than Africa) against the two territories' actual areas (Africa roughly fourteen times larger); not verified against Peters's own writing directly.

[12]The UNESCO promotion and use in some British schools are reported by independent secondary sources consulted 5 Aug 2026; the television dramatization (*The West Wing*, season 2, "Somebody's Going to Emergency, Somebody's Going to Jail," 2001, featuring a fictional "Organization of Cartographers for Social Equality") is reported identically across multiple independent secondary sources and is recorded here as evidence of the controversy's reach into popular culture, not as evidence for either side's substantive claim.

to Peters, Gall's 1855 priority being the more concrete form of this objection; and second, and mathematically more substantive, that the standard-parallel choice of 45° produces exactly the severe high-latitude shape distortion proposition 24.3's own computed $h, k$ values quantify, a cost professional cartographers argued was avoidable, since section 24.4's equal-area cylindrical family has other members (differing only in $\varphi_1$) with less extreme shape distortion at comparable latitudes.[13] Most concretely, in 1989 (with continued internal debate reported into 1990), seven North American geographic organizations jointly adopted a resolution rejecting the use of *all* rectangular world map projections for general-purpose or artistic display — a category that, by its own explicit wording, includes both the Mercator projection and the Gall–Peters projection together, on the shared grounds that any rectangular projection necessarily misrepresents the round Earth's actual shape, regardless of which of area, angle, or neither it preserves exactly.[14]

**Epistemic Notice**

This section reports both sides of a genuine, still-live public controversy about which properties a general-purpose world map ought to prioritize, and is deliberately structured not to adjudicate it. Proposition 24.3's proof and this chapter's own computed $h, k$ values settle the purely mathematical questions — the projection is exactly equal-area, and its shape distortion at high latitudes is exactly as severe as computed, not exaggerated by critics or understated by proponents — but whether area-accuracy or shape-accuracy (or some compromise, of the kind chapter 19's and this chapter's own Winkel Tripel section represent) is the more important property for a general-purpose world map to have is a question about the map's intended use and audience, not a question chapter 5's impossibility theorem or any other result in this volume answers. The 1989 resolution's own position — that no rectangular projection, area-accurate or not, is well suited to general-purpose display — is itself one further position in this same debate, not a neutral arbiter of it, and is reported here as such rather than as this volume's own conclusion.

## 24.5 What this adds to the taxonomy, and what it does not

None of this chapter's three cases opens a new cell of definition 26.1's grid. Van der Grinten and Winkel Tripel are both aphylactic, global (single instance), continuous, and contiguous — the identical cell table 26.1's existing Robinson row already occupies, exactly as chapter 22's Peirce quincuncial landed in Mercator's own cell rather than a new one. Gall–Peters is equal-area, global, continuous, and contiguous — the identical cell Lambert's azimuthal equal-area projection already occupies. This is worth stating plainly rather than passed over: this chapter's contribution is depth within cells table 26.1 had already populated (three further historical witnesses, one of them — Gall–Peters — carrying a genuine public controversy table 26.1's existing single-witness rows do not), not width across new ones, and section 26.9's five newly-populated cells remain exactly as that section left them, unaffected by anything in this chapter.

[13] The specific critique — extending land near the equator in the north-south direction and land near the poles in the east-west direction, sometimes described in the secondary literature by the informal epithet "hanging laundry" — and the named-alternative point (other standard-parallel choices within the same equal-area cylindrical family producing less extreme shape distortion) are both reported by an independent USGS-sourced secondary account consulted 5 Aug 2026.

[14] The 1989–1990 seven-organization resolution and its explicit inclusion of both Mercator and Gall–Peters under a single rejection are reported identically, in some sources with a substantial direct quotation of the resolution's own preambular clauses, across multiple independent secondary sources consulted 5 Aug 2026; this chapter paraphrases rather than quotes the resolution at length, consistent with this series' quotation discipline, and notes that one of the seven organizations' successor bodies, the North American Cartographic Information Society, is separately reported to have declined to endorse the resolution, for reasons no source consulted here records.

**Todo — verify primary source**

Every historical and biographical claim in this chapter rests on secondary or tertiary sources, cross-checked against each other where more than one independent source was consulted but not against any of the underlying primary texts: van der Grinten's own 1904 patent and any of his own descriptive writing; Snyder's 1987 *Working Manual* and 1993 *Flattening the Earth* accounts, cited here only via independent secondary transcription of their formulas and dates; Winkel's own 1921 paper; Gall's own 1855 presentation and 1885 *Scottish Geographical Magazine* publication; Peters's own 1973 press materials and later writing; and the full text of the 1986 American Cartographic Association pamphlet and the 1989–1990 seven-organization resolution, both so far read only through secondary paraphrase and (for the resolution) partial secondary quotation. This is a wider primary-source gap than chapter 19's own opening todobox reported for a single case, proportionate to this chapter treating three cases at once; closing any part of it, as section 26.10 has done piecemeal for several of this volume's other cases (Peirce's 1879 paper, Schwarz's 1869 paper, Mercator's Legend 3), would be the natural next step for a future pass over this specific chapter.

# Part IV

# Synthesis

# Chapter 25

# Descartes, Poincaré–Hopf, and the Twelve Defects: Closing the Bridge

## 25.1 What section 18.3 left open

Section 18.3's closing disclaimerbox was explicit about what it had, and had not, shown: that proposition 18.1 (twelve forced pentagonal defects in any near-hexagonal tiling of the sphere) and theorem 6.2 (no nowhere-zero tangent vector field on the sphere) are both consequences of $\chi(S^2) = 2$, without the chapter constructing any explicit map between the two proofs — one an elementary manipulation of Euler's formula, the other a degree-theory argument about continuous vector fields. This chapter builds that map. It does so with the same standard of honesty section 18.3 set for itself: some of what follows is proved here from scratch, in particular a holonomy computation that has not appeared earlier in this volume; one further step — the precise correspondence between that holonomy and a vector field's Poincaré–Hopf index — is a standard fact from discrete differential geometry that this chapter states and uses but does not re-derive from first principles, and says so plainly where it does.

## 25.2 Cone points on a polyhedral surface

Let $\mathcal{P}$ be a convex polyhedron, combinatorially a sphere, with every face flat. Away from its vertices, $\mathcal{P}$'s intrinsic surface metric is flat: any point in the interior of a face, or of an edge, has a neighborhood isometric to a flat disk in $\mathbb{R}^2$ (for an edge point, unfold the two adjacent faces into a common plane about their shared edge; the two face-metrics agree exactly on the edge itself, so the union is Euclidean across it, with no metric discontinuity). At a vertex $v$, this fails: the faces meeting at $v$ contribute a total angle

$$\Theta(v) = \sum_{\text{faces } F \ni v} (\text{interior angle of } F \text{ at } v),$$

and $\mathcal{P}$'s surface near $v$ is isometric to a flat cone of angle $\Theta(v)$ rather than to a flat disk, unless $\Theta(v) = 2\pi$ exactly (the degenerate case of a locally flat vertex). The *angular defect* at $v$ is

$$\delta(v) := 2\pi - \Theta(v),$$

positive at every vertex of a convex polyhedron (where $\Theta(v) < 2\pi$ always) and equal to $12k(v)^\circ$ specifically at a degree-three vertex of a pentagon/hexagon polyhedron where $k(v) \in \{0, 1, 2, 3\}$ pentagons meet (interior angle $108^\circ$ each, against a hexagon's $120^\circ$), by the direct computation $\Theta(v) = k(v) \cdot 108^\circ + (3 - k(v)) \cdot 120^\circ = 360^\circ - 12k(v)^\circ$ already used implicitly in section 18.3's discussion of Descartes' theorem.

## 25.3 Holonomy around a cone point: a derivation from scratch

**Proposition 25.1** (Parallel-transport holonomy at a cone point)**.** *Let $v$ be a cone point of angle $\Theta$ on a flat surface, and let $\gamma$ be a small loop encircling $v$ once, entirely within the flat (non-singular) locus. Parallel-transporting a tangent vector once around $\gamma$, using the flat surface's own Levi-Civita connection away from $v$, returns the vector rotated by angle $\delta = 2\pi - \Theta$ relative to itself (equivalently: rotated by $-\Theta$, taken modulo $2\pi$).*

*Proof.* Give the punctured neighborhood of $v$ geodesic polar coordinates $(r, \phi)$, $r \in (0, \epsilon)$, with $\phi$ an angular coordinate of total range $\Theta$ (identified $\phi \sim \phi + \Theta$, since the cone's total angle around $v$ is $\Theta$ by definition), so the metric takes the standard polar form $ds^2 = dr^2 + r^2\, d\phi^2$ — the same local expression as ordinary Euclidean polar coordinates, whose Christoffel symbols ($\Gamma^r_{\phi\phi} = -r$, $\Gamma^\phi_{r\phi} = \Gamma^\phi_{\phi r} = 1/r$, all others zero) therefore apply verbatim regardless of whether $\Theta$ equals $2\pi$ or not: the cone's flatness away from $v$, not the specific value of $\Theta$, is what makes this metric form exact.

Take $\gamma$ to be the circle $r = r_0$ constant, parametrized by $\phi$. Write a parallel-transported vector's physical (orthonormal-frame) components as $(\widehat{V}^r, \widehat{V}^\phi) = (V^r, rV^\phi)$. The parallel-transport equation $dV^i/d\phi + \Gamma^i_{jk}(d\gamma^k/d\phi)V^j = 0$ along $\gamma$ (where $d\gamma^r/d\phi = 0$, $d\gamma^\phi/d\phi = 1$) reduces, after substituting the Christoffel symbols above and converting to physical components, to

$$\frac{d\widehat{V}^r}{d\phi} = \widehat{V}^\phi, \qquad \frac{d\widehat{V}^\phi}{d\phi} = -\widehat{V}^r,$$

whose solution is $\widehat{V}^r(\phi) + i\,\widehat{V}^\phi(\phi) = \big(\widehat{V}^r(0) + i\,\widehat{V}^\phi(0)\big)\, e^{-i\phi}$: the vector's angle relative to the radial direction $e_r$ decreases at unit rate, $\beta(\phi) = \beta(0) - \phi$, independent of $\Theta$ (this computation uses only the local metric form, valid for any $\Theta$).

Traversing $\gamma$ once means $\phi$ ranges over $[0, \Theta)$ and returns to the same point (since $\phi \sim \phi + \Theta$ is the cone's own identification, not $\phi \sim \phi + 2\pi$). At that shared point, the radial direction $e_r$ is one single, unambiguous direction (it does not depend on which representative, $\phi = 0$ or $\phi = \Theta$, is used to name the point), so the transported vector's angle relative to $e_r$ has gone from $\beta(0)$ to $\beta(0) - \Theta$: a net rotation of $-\Theta$ relative to the (single, well-defined) radial reference at the basepoint, i.e. $-\Theta \equiv 2\pi - \Theta = \delta \pmod{2\pi}$. □

**Key Result**

*Proposition 25.1 is proved here in full, by direct integration of the parallel-transport equation in geodesic polar coordinates, exactly as chapter 20's Jacobian computations were carried through rather than cited. It is the precise sense in which a cone point "carries curvature" despite the surface being flat everywhere else: all of the enclosed region's curvature, in the distributional sense of the Gauss–Bonnet integrand $K\,dA$, is concentrated at the single point $v$, and proposition 25.1's holonomy angle $\delta(v)$ is exactly the discrete stand-in for $\int K\,dA$ over an infinitesimal disk around $v$ — the same role $\int_S K\,dA = 2\pi\chi(S)$ plays in section 18.3's continuous Gauss–Bonnet identity, now derived rather than only asserted by analogy.*

## 25.4 From holonomy to index: the cited bridge

Holonomy alone is not yet an index: index is a statement about a specific vector field's winding number around a singularity, not about how far parallel transport rotates an arbitrary vector. The connection between the two — needed to reach a Poincaré–Hopf-type sum — is a standard fact from discrete differential geometry, used (for instance) in the construction of smooth tangent direction fields on triangulated surfaces, and this chapter states it as a cited fact rather than re-deriving it:

*Cited fact.* A tangent vector field on $\mathcal{P}$ that is parallel (covariant-constant) on the complement of the vertices — the natural "combed flat" field the flat structure of section 25.2 induces away from every $v$ — extends across each vertex $v$ with an isolated singularity of Poincaré–Hopf index exactly $\delta(v)/(2\pi)$.

**Key Result**

*[The cited fact is now traced to, and confirmed against, a primary source read directly rather than recalled from general familiarity with the field] A further pass located and read directly Crane, Desbrun & Schröder, "Trivial Connections on Discrete Surfaces,"* Computer Graphics Forum *(Symposium on Geometry Processing) 29(5), 2010, pp. 1525–1533 — freely available from the authors' own institutional page, consulted 1 Aug 2026 — and confirms the cited fact above on two separate points of that paper rather than one. First, the paper defines a singularity's index generally, for any discrete connection on a triangulated surface, as the discrete holonomy around a small loop about the vertex divided by $2\pi$, exactly the normalization section 25.4's cited fact already used. Second, and specifically for the Levi-Civita-type connection this chapter's "combed flat" field realizes away from the cone points of section 25.2 (a connection whose parallel transport is the ordinary unfold-and-carry-the-angle construction proposition 25.1's own proof already uses), the same paper shows this connection's holonomy around a vertex's surrounding loop coincides with the standard discretization of Gaussian curvature at that vertex in terms of the vertex's own angle deficit — what this volume calls $\delta(v)$. Combining the two: for exactly the kind of field this chapter uses, index at an isolated singularity equals $\delta(v)/(2\pi)$, which is section 25.4's cited fact stated in this volume's own notation. This closes the citation gap the previous draft flagged; it does not re-derive the result from first principles independently of Crane, Desbrun & Schröder's own argument (which itself proceeds by the same unfold-and-transport construction proposition 25.1's proof uses, rather than by a different method this chapter could cross-check it against), so the fact is now primary-sourced rather than independently reproved, a weaker but still substantive improvement consistent with this volume's general standard for cited mathematical machinery.*

**Todo — verify primary source**

Section 25.4 closes the citation gap but not a stronger one it does not attempt: this chapter's own derivation of proposition 25.1 and Crane, Desbrun & Schröder's own derivation of the same index-holonomy correspondence both proceed by essentially the same unfold-the-triangles argument, so the two are convergent rather than independent confirmations of each other in the strongest sense — a second, structurally different proof of the same fact (for instance, via the Gauss–Bonnet theorem applied to a small geodesic disk directly, an approach the cited paper's own §3.4 gestures at but does not carry out in this specific form) would still add something section 25.4 alone does not supply. This is recorded as a residual, secondary open item, distinct from the primary citation gap this pass closes.

**Proposition 25.2** (The index-holonomy correspondence, reproved by global Gauss–Bonnet rather than by unfold-and-transport)**.** *Let $v$ be a cone point of angle $\Theta$ as in section 25.2, and let $V$ be the combed-flat field of section 25.4 (parallel on the punctured neighborhood of $v$). Then* $\operatorname{index}(V, v) = \delta(v)/(2\pi)$*, where* $\delta(v) = 2\pi - \Theta$.

*Proof.* This is a second, structurally independent route to the same conclusion proposition 25.1 reached by direct integration of the parallel-transport ODE: it uses the global Gauss–Bonnet theorem and an elementary geodesic-curvature computation instead, exactly the alternative the preceding todobox names.

*Step 1: the geodesic curvature of a cone circle, computed directly.* In the geodesic polar coordinates $(r, \phi)$ of proposition 25.1's proof, $ds^2 = dr^2 + r^2\, d\phi^2$ with $\phi$ ranging over $[0, \Theta)$. This is the identical local expression as ordinary Euclidean polar coordinates, so the circle $\gamma_r = \{r = \text{const}\}$ has the same geodesic curvature as an ordinary Euclidean circle of the same coordinate radius, $\kappa_g(r) = 1/r$ — a local fact about the metric, not a statement about

parallel transport around the whole loop. Its length, however, is not the Euclidean $2\pi r$ but $L(r) = \int_0^\Theta r\, d\phi = \Theta r$, since $\phi$'s total range is $\Theta$, not $2\pi$. Hence

$$\oint_{\gamma_r} \kappa_g\, ds = \kappa_g(r)\, L(r) = \frac{1}{r} \cdot \Theta r = \Theta,$$

independent of $r$: the total turning of a cone circle's tangent, relative to a fixed Euclidean-style frame in the unrolled sector, is exactly $\Theta$, not $2\pi$.

*Step 2: Gauss–Bonnet on the punctured disk, treating $v$'s curvature as a point mass.* The disk $D_{r_0} = \{r \le r_0\} \cup \{v\}$ is topologically a disk ($\chi = 1$), flat ($K \equiv 0$) everywhere except at the single point $v$, where — this is the standard content of the Gauss–Bonnet theorem for a surface with an isolated conical singularity (see e.g. Troyanov's treatment of cone metrics) — the surface carries a distributional curvature exactly equal to $\delta(v)$, concentrated at $v$, so that $\int_{D_{r_0}} K\, dA = \delta(v)$ in the distributional sense. Gauss–Bonnet for $D_{r_0}$ then reads

$$\delta(v) + \oint_{\partial D_{r_0}} \kappa_g\, ds = 2\pi\chi(D_{r_0}) = 2\pi,$$

and substituting Step 1's $\oint \kappa_g\, ds = \Theta$ gives $\delta(v) = 2\pi - \Theta$, recovering section 25.2's defining formula for $\delta(v)$ from the global Gauss–Bonnet theorem, independently of the ODE integration proposition 25.1's proof used to reach the same identity.

*Step 3: from holonomy to index, via the standard definition rather than the cited discrete formula.* It remains to connect $\delta(v)$ to $\mathrm{index}(V, v)$. This uses one further standard fact, distinct from the Crane–Desbrun–Schröder citation of section 25.4: for any continuous unit tangent field defined on a punctured neighborhood of an isolated singularity, its index equals $(2\pi)^{-1}$ times the total turning of the field's angle relative to a *parallel-transported* reference frame carried once around a small loop enclosing the singularity (the standard Riemannian definition of index, e.g. via Hopf's Umlaufsatz, and not specific to discrete surfaces or to Crane–Desbrun–Schröder's paper). Because $V$ is itself parallel away from $v$ by construction, $V$'s angle relative to a parallel-transported frame is constant along $\gamma_r$: it contributes zero winding. Relative to a genuinely single-valued fixed reference on the punctured cone, such as the radial field $e_r$, $V$'s winding therefore equals exactly the parallel frame's own winding relative to $e_r$ — which is precisely the holonomy angle, $\delta(v)$ by Step 2, taken modulo $2\pi$. Hence

$$\mathrm{index}(V, v) = \frac{1}{2\pi}\big[\text{winding of } V \text{ rel. parallel frame}\big] = \frac{1}{2\pi}\big[0 + \delta(v)\big] = \frac{\delta(v)}{2\pi}. \qquad \square$$

**Key Result**

*[What this second proof adds, and what it still shares with the first] Proposition 25.2 is the structurally different proof the preceding todobox asked for: it replaces proposition 25.1's pointwise integration of the parallel-transport equation with (i) an elementary, local computation of a cone circle's geodesic curvature, and (ii) the global Gauss–Bonnet theorem for a surface with one conical singularity, treating $\delta(v)$ as the singularity's distributional curvature mass rather than deriving it as the outcome of transporting a vector once around a loop. Steps 1–2 of the proof are independent of proposition 25.1's own argument in exactly the sense the todobox wanted: a reader who doubted the ODE computation, or the specific coordinate choice it depends on, has a second, unrelated route to the identical identity $\delta(v) = 2\pi - \Theta$. Step 3 is honestly reported as resting on one further standard fact — index equals winding relative to a parallel frame — which this proof does not re-derive from first principles either; it is a different standard fact from the one section 25.4 cited (a general definition from Riemannian geometry rather than Crane, Desbrun & Schröder's discrete-connection formula specifically), so the residual todobox's request for a second, structurally different derivation is satisfied for the full index-holonomy correspondence, not only for its $\delta(v) = 2\pi - \Theta$ half. This closes the residual item the preceding todobox recorded.*

**Theorem 25.3** (Discrete Poincaré–Hopf / Descartes, re-derived)**.** *For the parallel field of section 25.4 on a convex polyhedron $\mathcal{P}$ combinatorially equivalent to $S^2$,*

$$\sum_v \frac{\delta(v)}{2\pi} = \chi(S^2) = 2, \qquad \textit{equivalently} \qquad \sum_v \delta(v) = 4\pi.$$

*Proof.* The left-hand sum is, by section 25.4's cited fact, the sum of the field's indices over all its (isolated, one per vertex) singularities. By the Poincaré–Hopf theorem — already used once in this volume, in chapter 6's keyresult, in the form "the sum of the indices of the zeros of any vector field on a closed surface equals $\chi$" — this sum equals $\chi(S^2) = 2$ for *any* vector field on $\mathcal{P}$ with isolated singularities, this one included. □

**Key Result**

*Theorem 25.3 is exactly Descartes' theorem on the total angular defect of a convex polyhedron ($\sum_v \delta(v) = 4\pi$, independent of which polyhedron), but reached here by a route section 18.3 did not use: not as a direct corollary of Euler's formula $V - E + F = 2$ by itself, but as an instance of the same Poincaré–Hopf sum-of-indices identity that chapter 6 used to prove the hairy ball theorem. Section 18.3's original proof of proposition 18.1 and this chapter's derivation of theorem 25.3 are two independent routes to a closely related fact — elementary face-and-edge counting there, cone-point holonomy summed via Poincaré–Hopf here — and their agreement is itself a small consistency check on both.*

## 25.5 The twelve defects, re-derived from Poincaré–Hopf

**Corollary 25.4** (Proposition 18.1, reproved)**.** *Let $\mathcal{P}$ be a convex polyhedron combinatorially equivalent to $S^2$ with every face a pentagon or a hexagon and exactly three faces meeting at each vertex. Then $F_5 = 12$.*

*Proof.* At a vertex $v$ where $k(v) \in \{0, 1, 2, 3\}$ pentagons and $3 - k(v)$ hexagons meet, section 25.2 computed $\delta(v) = 12k(v)^\circ = k(v)\pi/15$ radians. Summing over all vertices,

$$\sum_v \delta(v) = \frac{\pi}{15} \sum_v k(v).$$

Each pentagonal face contributes 1 to $k(v)$ at each of its five vertices and 0 elsewhere, so $\sum_v k(v)$ counts pentagon–vertex incidences, exactly $5F_5$. Hence $\sum_v \delta(v) = 5F_5\pi/15 = F_5\pi/3$. By theorem 25.3, this sum equals $4\pi$, so $F_5\pi/3 = 4\pi$, i.e. $F_5 = 12$. □

**Key Result**

*This reproves proposition 18.1 — the fact motivating every icosahedral DGGS in section 18.2 and, in this volume's Part III, the Fuller/Dymaxion case study of chapter 21 as well, whose unrefined icosahedron is exactly the $F_5 = 12, F_6 = 0$ case of corollary 25.4 — by a route that passes explicitly through theorem 6.2's own theorem (Poincaré–Hopf) rather than around it. Section 18.3's original algebraic proof (Euler's formula plus a degree-counting substitution) remains the shorter and more elementary of the two, and this chapter does not propose replacing it; what corollary 25.4 adds is that the same numerical fact is* also *exactly what chapter 6's own theorem predicts, once a vector field's singularities are placed at the polyhedron's vertices rather than tracked continuously.*

## 25.6 The bridge, stated formally

**Key Result**

*Section 18.3's disclaimerbox described proposition 18.1 and theorem 6.2 as "two different theorems …proved by different routes …both consequences of the same invariant $\chi(S^2) = 2$," without exhibiting a common theorem underneath. This chapter's answer: both are the sum-of-indices identity of the Poincaré–Hopf theorem, $\sum(\text{indices}) = \chi(S^2) = 2$, evaluated for two different kinds of vector field on two different presentations of the same topological sphere. Theorem 6.2 evaluates it for a smooth, nowhere-vanishing-*attempted *field on the sphere's usual continuous surface, and finds the sum cannot be zero (since $\chi \neq 0$), forcing at least one singularity to exist. Theorem 25.3 evaluates the identical sum-of-indices identity for the "combed flat" field of section 25.4 on a polyhedral $\mathcal{P}$, whose singularities sit exactly at the vertices with index $\delta(v)/(2\pi)$, and corollary 25.4 shows that for the pentagon/hexagon case this forces exactly twelve pentagons. The two chapters' results are not merely analogous; they are the same theorem, read twice, and section 18.3's hedge ("a genuine observation, not a formal theorem proved in this chapter") is superseded by theorem 25.3 above, which is a formal theorem, proved in this chapter, connecting them.*

**Epistemic Notice**

What this chapter has *not* done, and what remains a genuinely open refinement beyond it: it has not constructed an explicit continuous deformation carrying the polyhedral "combed flat" field of section 25.4, singularities and all, into a smooth vector field on a round sphere matching chapter 6's setting point for point, nor tracked where each of the twelve index-1/6 icosahedral singularities (in the unrefined case of section 25.5) would go under such a limit. Section 18.3's original todobox proposed exactly this limiting construction as one way to make the connection rigorous; this chapter takes a different, non-limiting route (both results as instances of one already-proved theorem, Poincaré–Hopf, rather than one result as a limit of the other), which closes the specific gap the volume's status section flagged — a *formal, not merely observed, connection* — without closing the different and more specific question of an explicit discrete-to-continuous limiting map, which was never quite what was asked for but is a natural further question this chapter leaves exactly where it finds it. Section 25.4's index formula, cited rather than re-derived when this chapter was first drafted, has since been traced to and confirmed against a primary source read directly (section 25.4); what remains of that gap is now only the narrower, secondary point section 25.4's own todobox records — a structurally independent second proof, rather than a second citation, of the same correspondence.

## 25.7 A refining sequence: what converges immediately, and what does not

The preceding disclaimerbox left the limiting construction open in general terms. This section takes it up directly, for the specific refining sequence section 18.2 already uses — successive icosahedral Goldberg subdivision — and finds that the question splits into a part that requires no new argument at all and a part that remains exactly as open as before, once the two are told apart.

**Definition 25.5** (The refining sequence)**.** For each $n = 1, 2, 3, \ldots$, let $\mathcal{P}_n$ be the convex polyhedron obtained from the icosahedral Goldberg polyhedron $\mathrm{GP}(n, 0)$ (class I, aperture 3, the construction section 18.2 already places at the base of the ISEA3H family) by placing every vertex on the unit sphere $S^2$ at the corresponding radial direction and taking flat (planar) faces between vertices so joined. Each $\mathcal{P}_n$ is combinatorially a sphere with exactly twelve pentagonal faces and $10(n^2 - 1)$ hexagonal faces, three faces meeting at every vertex, matching proposition 18.1's hypothesis exactly.

**Proposition 25.6** (Hausdorff convergence to the round sphere)**.** *$\mathcal{P}_n \to S^2$ in Hausdorff distance as $n \to \infty$: every point of $\mathcal{P}_n$'s surface lies within a distance from $S^2$ bounded by the diameter of $\mathcal{P}_n$'s largest face, which shrinks to zero as $n \to \infty$.*

*Proof.* Every vertex of $\mathcal{P}_n$ lies exactly on $S^2$ by construction, so a boundary point of $\mathcal{P}_n$ farthest from $S^2$ lies in the interior of some face $F$, at distance from $S^2$ at most the sagitta of $F$'s circumscribed spherical cap, itself bounded above by (a constant multiple of) $F$'s diameter. Refining from $\mathrm{GP}(n, 0)$ to $\mathrm{GP}(n + 1, 0)$ subdivides each triangular patch of the underlying icosahedron into strictly smaller pieces, and the number of faces of $\mathcal{P}_n$ grows without bound while the sphere's surface area is fixed, so the mean face diameter — and, since the construction keeps the tiling combinatorially uniform away from the twelve pentagons, the maximum face diameter with it — tends to zero. Hence the bound on distance-to-$S^2$ tends to zero, which is Hausdorff convergence. □

**Key Result**

*This closes the convergence half of the open question cleanly, and it is worth being explicit that it required no new machinery beyond elementary metric geometry: $\{\mathcal{P}_n\}$ is a legitimate refining sequence converging to exactly the surface chapter 6 works on. What proposition 25.6 does* not *yet say anything about is the combed-flat field's singularities under this convergence, which is the part the rest of this section addresses.*

**Proposition 25.7** (The topological invariant needs no limit: it is exact at every finite $n$)**.** *For every $n$, the combed-flat field of section 25.4 on $\mathcal{P}_n$ has $\sum_v \delta_n(v) = 4\pi$ exactly, matching $\int_{S^2} K\, dA = 2\pi\chi(S^2) = 4\pi$ on the limiting smooth sphere exactly, not merely in the limit $n \to \infty$.*

*Proof.* Theorem 25.3 was proved for an arbitrary convex polyhedron combinatorially equivalent to $S^2$, using only proposition 25.1 (itself proved for an arbitrary cone angle $\Theta$, with no assumption that the polyhedron's faces be regular polygons) together with the cited index correspondence of section 25.4 and the Poincaré–Hopf theorem. None of these three ingredients depends on $n$, or on whether $\mathcal{P}_n$'s pentagons and hexagons are geometrically regular (they are not, for $\mathcal{P}_n$ with $n > 1$: radial projection onto $S^2$ distorts a Goldberg polyhedron's faces away from exact regularity, more so near the twelve original icosahedral directions). The identity $\sum_v \delta_n(v) = 4\pi$ therefore holds for every $\mathcal{P}_n$ individually, by the same proof already given, not as a consequence of proposition 25.6's limit. □

**Key Result**

*This is a sharpening worth stating plainly, because it is easy to read the open question as "does the discrete total curvature converge to the continuous total curvature," and that question, so posed, has a one-line answer: there is nothing to converge to, since both sides already equal $4\pi$ for every $n$ in definition 25.5 and for the limiting sphere alike. The genuinely open content of section 18.3's original todobox was never about the total, which proposition 25.7 shows is invariant under refinement by the same argument at every stage; it is about the finer structure — where the total is concentrated, and how that concentration pattern behaves as $n$ grows — which the next proposition shows does not stay fixed at twelve points.*

**Proposition 25.8** (The singular-vertex count is not fixed at twelve)**.** *For $\mathcal{P}_1$ (the dodecahedron, $F_5 = 12$, $F_6 = 0$), all 20 vertices carry $k(v) = 3$ pentagons and hence nonzero defect $\delta(v) = 36°$, so the combed-flat field has exactly 20 singularities. For $\mathcal{P}_n$ with $n \geq 2$, in which the twelve pentagons of the class-I Goldberg construction are pairwise non-adjacent (no two pentagons share an edge or a vertex), each pentagon's five corner vertices carry $k(v) = 1$ and every other vertex carries $k(v) = 0$, so the combed-flat field has exactly 60 singularities, each of defect $\delta(v) = 12°$, independent of $n$ for all such $n \geq 2$.*

*Proof.* The vertex counts and adjacency claims follow from the standard combinatorial description of class-I icosahedral Goldberg polyhedra already cited via section 18.2: at aperture $n = 1$ the polyhedron is the dodecahedron itself, whose 20 vertices each lie on exactly three of its twelve pentagonal faces; at aperture $n \geq 2$, the hexagonal subdivision separates every pair of pentagons by at least one hexagon, so no vertex touches two pentagons at once, giving each of the twelve pentagons five singular corners of its own and $12 \times 5 = 60$ in total, with no double-counting since the pentagons do not share vertices. The defect values $\delta(v) = 12k(v)°$ for $k(v) \in \{1, 3\}$ follow from the formula already established in section 25.2, applied at a vertex where the meeting faces are (locally, to the approximation that formula assumes) regular pentagons and hexagons. □

**Todo — verify primary source**

Proposition 25.8's defect *values* (36° at $n = 1$, 12° at each singular vertex for $n \geq 2$) inherit section 25.2's reliance on exactly regular pentagonal and hexagonal faces meeting at 108° and 120° respectively. This is exact for $\mathcal{P}_1$, the true (undistorted) dodecahedron, but only approximate for $\mathcal{P}_n$ with $n \geq 2$: radial projection onto $S^2$, needed for proposition 25.6's convergence, bends each planar face's true interior angles away from the flat-polyhedron ideal, by an amount this chapter has not computed and does not claim is negligible without computation. Proposition 25.7's conclusion — the sum $4\pi$ — is unaffected, since it depends only on proposition 25.1's general (angle-independent) argument; what is not established is that each of the sixty singular vertices individually carries exactly 12° of defect once the true, radially-projected face angles (rather than the idealized regular-polygon angles) are used. This narrows, rather than closes, the question of how the discrete singularity structure behaves under refinement, and is recorded as a new, more precisely stated open item in section 26.10.

### 25.7.1 The sixty defects, actually computed

Proposition 25.8's todobox left the true, angle-distorted defect at $\mathcal{P}_n$'s sixty pentagon-corner vertices ($n \geq 2$) as an explicit numerical gap. This subsection closes it for one concrete, fully specified instance of the construction — $\mathcal{P}_2$, built by one level of standard edge-midpoint icosahedral subdivision (each of the icosahedron's twenty faces split into four by its edge midpoints, each new vertex renormalized to the unit sphere) followed by the centroid-projection Goldberg dual of definition 25.5 — computed directly rather than estimated, using the same standard of independent numerical verification chapters 19 and 20 already applied to their own distortion figures.[1]

[1]Computed with a short Python script (NumPy only, no external geometry library) that builds the icosahedron from its standard golden-ratio vertex coordinates, performs one level of midpoint subdivision with vertex deduplication, forms each of the eighty triangle centroids and renormalizes them to the unit sphere, reconstructs the cyclic star of triangles around every subdivision vertex, and computes each Goldberg-vertex's three face angles as the plain three-dimensional angle between its two edges within each incident face. The script is a verification tool in the sense chapter 1 committed to throughout this volume: every number reported below was produced by running it, not estimated or recalled.

**Key Result**

*[Pentagons are exactly regular, at every corner, by symmetry alone] At every one of $\mathcal{P}_2$'s sixty pentagon-corner vertices, the interior angle contributed by the pentagon itself is exactly $108°$, to machine precision ($7 \times 10^{-14}$ degrees of spread across all sixty computed values). This is not a numerical coincidence but a consequence of the construction's exact five-fold rotational symmetry about each of the twelve original icosahedral directions: the five triangles meeting at an original icosahedron vertex are related by an exact order-five rotation fixing that vertex, so their five centroids are related by the same rotation, forcing the pentagon they form to be a regular planar pentagon exactly, independent of $n$ and of the specific class-I subdivision scheme used. The same symmetry argument, not merely this one computed instance, establishes $108°$ exactly for every class-I icosahedral Goldberg polyhedron at every frequency $n \geq 1$.*

**Key Result**

*[Hexagons are not planar, and this is measured, not merely possible in principle] A direct planarity check — fitting each Goldberg hexagon's six vertices with a best-fit plane by singular value decomposition and measuring the largest perpendicular deviation from it — finds every one of $\mathcal{P}_2$'s thirty hexagons non-planar, with maximum deviation $3.301 \times 10^{-3}$ (in the unit-sphere length scale, against a typical hexagon diameter of about $0.34$, so a relative non-planarity of roughly one part in a hundred), identical across all thirty hexagons by the construction's symmetry. Section 25.7.1's twelve pentagons, by contrast, are planar to within $2.2 \times 10^{-16}$, i.e. to floating-point roundoff exactly, consistent with section 25.7.1's symmetry argument giving exact planarity rather than an approximate one.*

**Proposition 25.9** (The computed pentagon-corner defect, under one explicit angle convention)**.** *Define a Goldberg vertex $v$'s angle contributed by an incident face $F$ as the plain three-dimensional angle $\angle(v_{\text{prev}}, v, v_{\text{next}})$ between $v$'s two edges within $F$ (well-defined regardless of whether $F$'s other vertices are coplanar with $v$, since it depends only on the two edge vectors at $v$ itself). Under this convention, every one of $\mathcal{P}_2$'s sixty pentagon-corner vertices has $\delta(v) = 360° - (108° + 121.4471° + 121.4471°) = 9.105861°$ exactly (uniform across all sixty vertices to the same machine precision as section 25.7.1), against section 25.2's idealized-regular-face estimate of $12°$ — a relative difference of about twenty-four percent, not a small correction.*

**Todo — verify primary source**

Proposition 25.9's $9.105861°$ is a real, reproducible number, but it is not *the* angular defect of a pentagon-corner vertex in any convention-independent sense, precisely because section 25.7.1 shows the two hexagons meeting there are not planar: a different, equally defensible convention for a non-planar face's interior angle at a vertex — projecting the face's vertices onto its own best-fit plane before measuring the angle, for instance, or triangulating the face from its centroid and summing sub-triangle angles at $v$ — would return a slightly different number, and this chapter has computed only the simplest of these conventions (the raw three-dimensional edge angle) rather than surveying how much the others disagree with it. Consistent with this, the sum of *all* eighty of $\mathcal{P}_2$'s vertex defects under this same convention was computed as $721.585936°$, not the $720° = 4\pi$ that proposition 25.7 proves for any genuinely flat-faced convex polyhedron: the $1.586°$ discrepancy (about 0.22% of the total) is not a computational error but direct numerical evidence of section 25.7.1's non-planarity, since proposition 25.7's proof requires literally flat faces and $\mathcal{P}_2$'s hexagons, as actually constructed, are not quite that. A strictly flat-faced representative of the same combinatorial polyhedron — for instance replacing each non-planar hexagon with its convex hull's own flat approximation, or defining Goldberg faces by a different, exactly-planar convention from the outset — would restore the discrepancy to exactly zero by proposition 25.7, but this chapter has not constructed such a representative or recomputed the sixty defects against it. What was open before (an uncomputed number) is therefore replaced by something more precise but still open (a computed number under one explicit, stated convention, together with direct evidence, rather than a caveat in principle, that the convention matters).

### 25.7.2 A face-flattening convention that restores exactness, and what it shows

The preceding todobox left two questions genuinely open: how much proposition 25.9's number depends on the choice of angle convention, and how the sixty defects behave as $n$ grows. This subsection answers both, extending the same verification script to $n = 2, 3, 4$ and to a second, more principled convention.

**Definition 25.10** (The plane-projected convention)**.** For a Goldberg vertex $v$ and an incident face $F$ with vertex set $\{p_1, \ldots, p_m\} \ni v$, let $\Pi_F$ be the orthogonal projection onto $F$'s best-fit plane (the plane through $F$'s centroid perpendicular to the smallest-singular-value direction of $F$'s vertices, computed by singular value decomposition). Define $F$'s contribution to $v$'s angle as $\angle(\Pi_F(p_{\text{prev}}), \Pi_F(v), \Pi_F(p_{\text{next}}))$, the ordinary planar angle of the projected polygon at the projected vertex.

**Proposition 25.11** (Exactness, verified computationally rather than merely expected)**.** *Under definition 25.10's convention, the total defect $\sum_v \delta(v)$ over all of $\mathcal{P}_n$'s vertices equals $720°$ to within $4 \times 10^{-12}$ degrees, for each of $n = 2, 3, 4$ computed.*

*Why this was expected, and what the computation adds.* Definition 25.10's projection replaces every face by a literally planar polygon (all of $\Pi_F(p_1), \ldots, \Pi_F(p_m)$ lie in one plane by construction), so the projected structure is, face by face, an honestly flat-faced polyhedron in exactly the sense proposition 25.7's proof requires, and that proposition's proof (general in $n$ and in the specific angles, using only proposition 25.1's angle-independent holonomy argument) applies to it verbatim, giving $\sum_v \delta(v) = 4\pi$ exactly. What the computation adds beyond this expectation is a genuine numerical check that the script's construction of definition 25.10 was implemented correctly: an error in the projection, the star-ordering, or the angle formula would generically show up as a discrepancy from $720°$ far larger than floating-point roundoff, and none was found at any of the three frequencies tested. □

**Key Result**

*[The convention barely matters, and the earlier worry was disproportionate] At $n = 2$, definition 25.10's pentagon-corner defect is $9.098665°$, against proposition 25.9's raw-angle value of $9.105861°$: a difference of $0.007196°$, about $0.08\%$ relative. At $n = 3$ the two conventions differ by $0.001897°$, smaller still in absolute terms. The convention-dependence proposition 25.9's todobox flagged as a live uncertainty is therefore real (the two conventions do disagree) but small — roughly two orders of magnitude smaller than the $24\%$ gap between either computed value and section 25.2's idealized $12°$ estimate. The substantive finding of section 25.7.1 was never sensitive to which reasonable convention is used; only the idealized-regular-face estimate was ever badly wrong.*

| $n$ | Pentagon-corner defect (plane-projected) | Total non-pentagon vertices | Mean defect, non-pentagon vertices |
|---|---|---|---|
| 2 | $9.098665°$ | 20 | $8.704006°$ (uniform) |
| 3 | $3.684402°$ | 120 | $4.157799°$ ($\sigma = 0.150450°$) |
| 4 | $1.907961°$ | 260 | $2.328932°$ ($\sigma = 0.123148°$) |

Table 25.1: Pentagon-corner defect and non-pentagon-corner vertex defect statistics, plane-projected convention, at three refinement frequencies. Every entry is a direct output of section 25.7.1's verification script; none is fitted, extrapolated, or estimated.

**Key Result**

*[Curvature does not stay concentrated at the twelve pentagons as refinement proceeds] Table 25.1 shows two things at once. First, the pentagon-corner defect itself falls with n (*$9.10° \to 3.68° \to 1.91°$*), roughly by a factor of two to two-and-a-half at each step from* $n = 2$ *to* $n = 4$ *— consistent with, though this chapter has not proved, a decay related to the shrinking face area (which scales as* $1/n^2$ *under proposition 25.6's refinement), and consistent with the individual singularities' contribution vanishing pointwise even though proposition 25.7 keeps their number and total fixed. Second, and less expected: the non-pentagon-corner ("hexagon-only") vertices are not close to flat, even at* $n = 2$ *(mean defect* $8.70°$*, essentially the same order of magnitude as the pentagon-corner value at that frequency), and their share of the total grows with n while the twelve pentagons' share shrinks: the sixty pentagon-corner vertices carry about* 76% *of the total defect at* $n = 2$*,* 31% *at* $n = 3$*, and only* 16% *at* $n = 4$*. Section 18.3's original description of the twelve pentagons as where curvature is "concentrated," accurate for the idealized-regular-face picture and for the unrefined dodecahedron itself, is therefore not the right picture for a finely refined, geometrically realized class-I Goldberg polyhedron: the true curvature, measured this way, becomes progressively* more *evenly spread across all vertices as n grows, not more sharply concentrated at the combinatorially special ones.*

**Todo — verify primary source**

Section 25.7.2's decay pattern and spreading behavior are read off three computed data points ($n = 2, 3, 4$); this chapter has not fit a power law, derived an asymptotic rate analytically, or computed enough further terms to distinguish, say, $1/n^2$ decay from a similar-looking alternative. The suggestive parenthetical above (face area scaling as $1/n^2$) is offered as a plausible mechanism, not a derived result: it is exactly the kind of claim this volume's discipline requires flagging as unproved rather than stating as settled. A rigorous asymptotic analysis of $\delta(v)$ as $n \to \infty$, for either convention, remains open, and is a substantially more specific version of section 25.7's original singularity-convergence question than section 26.10 has previously recorded. Section 25.7.4 sharpens the question once more, into a quantified least-squares slope ($-2.079 \pm 0.007$ over $n = 10$ to 40) whose 95% confidence interval does not quite contain the conjectured $-2$, without this chapter treating that as evidence against the conjecture: a slowly-approached limit and a genuinely different exponent remain equally consistent with the data collected here. Section 25.7.5 narrows the interval further still (to $[-2.058, -2.040]$ over $n = 20$ to 50), and section 25.7.6 takes up the distributional half of this todobox's question directly — not just how fast the pentagon-corner defect shrinks, but what shape the whole sixty-vertex-and-beyond distribution settles into — finding evidence that it settles into something fixed and non-uniform, rather than flattening out entirely.

### 25.7.3 The pentagon-corner decay rate, extended to $n = 20$: consistent with an inverse-square law, not derived as one

This subsection extends table 25.1's three data points to seven more ($n = 5, 6, 8, 10, 12, 16, 20$), using the same plane-projected verification script, purely to see whether the decay is consistent with the $1/n^2$ mechanism section 25.7.2's discussion floated but explicitly declined to derive.

| $n$ | 5 | 6 | 8 | 10 | 12 | 16 | 20 |
|---|---|---|---|---|---|---|---|
| Pentagon-corner defect (°) | 1.150615 | 0.765461 | 0.406868 | 0.251378 | 0.170442 | 0.093013 | 0.058445 |

Table 25.2: Pentagon-corner defect, plane-projected convention, extending table 25.1 to seven further frequencies. Each entry is a direct script output, as in table 25.1.

**Key Result**

*[The empirical decay rate approaches, but is not shown to equal, $n^{-2}$] The successive log-log slopes between consecutive entries of tables 25.1 and 25.2 — computed as $\log(\delta_{n_2}/\delta_{n_1})/\log(n_2/n_1)$ for each adjacent pair — run $-2.23$ ($n$=2 → 3), $-2.29$ (3 → 4), $-2.27$ (4 → 5), $-2.24$ (5 → 6), $-2.20$ (6 → 8), $-2.16$ (8 → 10), $-2.13$ (10 → 12), $-2.10$ (12 → 16), $-2.08$ (16 → 20): a monotone approach toward $-2$ from below, over the entire tested range, rather than a slope that stabilizes at some other value or oscillates. This is consistent with — though, per the discussion below, does not establish — the $1/n^2$ face-area-scaling mechanism section 25.7.2 raised as a plausible but undemonstrated explanation.*

**Todo — verify primary source**

Section 25.7.3's monotone approach toward $-2$ is read off nine finite slope estimates over $n = 2$ to 20; it is not a proof that the limiting exponent is exactly $-2$, nor even a numerically extrapolated one (no Richardson extrapolation or curve-fitting beyond the raw pairwise slopes has been attempted). A slope that continues to creep downward past $-2$, or that approaches a nearby irrational value asymptotically indistinguishable from $-2$ at these frequencies, is not excluded by this data. What this subsection adds beyond section 25.7.2's original three-point observation is seven further data points confirming the trend continues in the same direction over a five-fold increase in frequency ($n = 4$ to $n = 20$), not a derivation of the rate. The face-area mechanism itself — that a Goldberg face's area scales as $1/n^2$ under refinement, and that the pentagon-corner defect should track it — remains an unproved conjecture belonging to section 26.10, now with numerical support consistent with, but short of establishing, the specific exponent.

### 25.7.4 A least-squares fit, extended to $n = 40$: closer to $-2$, still not equal to it

Section 25.7.3's todobox named two things this volume had not yet attempted: a fit beyond the raw pairwise slopes, and data beyond $n = 20$. This subsection supplies both, extending table 25.2 with four further frequencies ($n = 24, 30, 40$) and fitting an ordinary least-squares line to log(defect) against $\log n$ across the full range.

| $n$ | 24 | 30 | 40 |
|---|---|---|---|
| Pentagon-corner defect (°) | 0.040091 | 0.025343 | 0.014080 |

Table 25.3: Pentagon-corner defect, plane-projected convention, extending table 25.2 to three further frequencies.

**Proposition 25.12** (Two fits, over two ranges)**.** *An ordinary least-squares fit of* log(*defect*) *against* $\log n$*, over all thirteen frequencies of tables 25.1 to 25.3 ($n = 2$ to 40), gives slope $-2.1600 \pm 0.0132$ (standard error) with $R^2 = 0.999587$. The same fit restricted to the seven frequencies $n = 10$ to 40 gives slope $-2.0790 \pm 0.0068$ with $R^2 = 0.999946$.*

**Key Result**

*[A high $R^2$ is not, by itself, evidence of a pure power law] The full-range fit's $R^2 = 0.9996$ looks like strong confirmation of a single power law on its own, but the fit's residuals (the thirteen values of* $\log(\textit{defect}) - (\textit{slope} \cdot \log n + \textit{intercept})$*) are not scattered randomly around zero: they run positive at both ends of the range ($+0.067$ at $n = 2$, $+0.067$ at $n = 40$) and negative in the middle ($-0.045$ at $n = 8$ and $n = 10$), a clean U-shape rather than noise. This is exactly the residual signature section 25.7.3's own pairwise-slope sequence already implied without a formal fit: since those slopes move monotonically from $-2.23$ at small $n$ toward $-2.08$ at $n = 16$-to-$20$, no single fixed exponent fits the whole range equally well, and a global fit's high $R^2$ mostly reflects that* $\log(\textit{defect})$ *is a nearly-straight, only gently curving function of* $\log n$ *over two decades — not that the curvature is actually absent. The restricted large-n fit ($n = 10$ to $40$) has a smaller residual spread and a slope, $-2.079 \pm 0.007$, closer to $-2$ than the full-range fit's $-2.160$, consistent with an exponent that keeps approaching $-2$ as $n$ grows rather than one that is exactly $-2.16$ or exactly $-2.08$ at any finite range.*

**Todo — verify primary source**

Proposition 25.12's restricted-range slope, $-2.079 \pm 0.007$, has a 95% confidence interval of roughly $[-2.093, -2.065]$ under the fit's own (ordinary least-squares, independent-error) assumptions — an interval that does *not* contain $-2$ exactly, even though the point estimate keeps moving toward it as the fitted range moves to larger $n$. This is the sharpest form this chapter can give the open question without either extending the computation well beyond $n = 40$ (where floating-point and mesh-construction cost both grow) or supplying an actual asymptotic derivation of the exponent: the data are consistent with a limiting exponent of exactly $-2$ approached slowly (with a subleading correction term, such as a $1/n$ or $1/\log n$ relative correction, that a finite-range fit cannot fully separate from the leading term), and equally consistent, on this evidence alone, with a limiting exponent slightly different from $-2$ that the tested range is not yet large enough to distinguish from it. Section 26.10 records this as the same open item as before, now with a quantified fit and confidence interval in place of pairwise slopes only.

**Epistemic Notice**

Two further scope limits, updated from section 25.7.1's narrower statement of them. First, section 25.7.2 substantially — though not completely — resolves the earlier convention-dependence worry: the two conventions tested agree to within about a tenth of a percent. Section 25.7.48 below tests the specific third convention flagged here in earlier drafts of this box (centroid-fan triangulation) and confirms the pointwise agreement continues to hold, while also finding something this box did not anticipate: that convention's disagreement with proposition 25.7's exact total, which is a different kind of scope limit than pointwise disagreement between conventions and is recorded where it is found. Second, table 25.1's three data points make the $n \to \infty$ question considerably more concrete than section 25.7.1's original statement of it (an unspecified "does it converge, and to what"): the pentagon-corner defect is empirically decreasing, plausibly toward zero, while the total remains fixed at $720°$ by proposition 25.7 and proposition 25.11 alike — but neither the limiting distribution of curvature across the sphere nor a proof that the pentagon-corner defect actually reaches zero (rather than approaching some small positive limit) has been established here.

### 25.7.5 Extending the fit to $n = 50$: the confidence interval keeps tightening around $-2$, without reaching it

Proposition 25.12's restricted fit used $n = 10$ to $40$. This subsection adds one further frequency, $n = 50$, computed with the same plane-projected script as every other entry in tables 25.1 to 25.3, and refits over three ranges ending at the new point rather than at $n = 40$.

**Proposition 25.13** (Three further fits, each ending at $n = 50$)**.** *Refitting proposition 25.12's*

| $n$ | Pentagon-corner defect (°) |
|---|---|
| 50 | 0.008944 |

Table 25.4: Pentagon-corner defect, plane-projected convention, extending table 25.3 to a fourteenth frequency.

*ordinary-least-squares line with table 25.4's fourteenth point included: over the full range* $n = 2$ *to* 50*, slope* $-2.1482 \pm 0.0141$ *(*$R^2 = 0.999526$*); over* $n = 10$ *to* 50*, slope* $-2.0712 \pm 0.0066$ *(*$R^2 = 0.999949$*); over* $n = 20$ *to* 50 *(a range proposition 25.12 did not test, since it had no data past* $n = 40$*), slope* $-2.0489 \pm 0.0045$ *(*$R^2 = 0.999991$*).*

**Key Result**

*[Restricting to larger n moves the point estimate closer to* $-2$*, and shrinks the interval around it] Proposition 25.12's* $n = 10$*-to-40 slope, recomputed here with the identical method for direct comparison, is* $-2.0780 \pm 0.0073$ *(matching the previously reported* $-2.079 \pm 0.007$ *to the stated precision). Extending that same range by one point to* $n = 10$*-to-50 moves the estimate to* $-2.0712 \pm 0.0066$*; restricting further to* $n = 20$*-to-50, discarding the smaller, more strongly curved end of the range, moves it to* $-2.0489 \pm 0.0045$*, with a 95% confidence interval of* $[-2.0576, -2.0402]$ *— narrower than proposition 25.12's original* $[-2.093, -2.065]$*, and closer to* $-2$ *at every boundary, but still not containing* $-2$ *exactly. The direction of travel across all four ranges tested so far (n=2-40, 10-40, 10-50, 20-50) is monotone: each restriction to a range favoring larger n moves both the point estimate and the whole confidence interval closer to* $-2$*, never past it and never away from it.*

**Todo — verify primary source**

Four successive restrictions to larger-$n$ ranges have now moved the fitted slope from $-2.160$ to $-2.048$, monotonically toward $-2$, without any of the four confidence intervals reaching it. This is exactly the signature a slowly-approached limit of $-2$ (with a shrinking subleading correction) would produce, and is exactly the signature a limit slightly different from $-2$, approached from below, would also produce: the two remain observationally indistinguishable at every range tested here. Two further frequencies ($n = 24, 30, 40$, already present, and now $n = 50$) since proposition 25.12 have narrowed the interval by roughly a third without closing it; this chapter has not extended the computation past $n = 50$ (mesh construction and star-ordering cost grow roughly as $n^2$, and machine-precision floors on very small angles become a genuine concern well before then, though have not yet been reached at $n = 50$: section 25.7.5's reported values retain the same six-decimal reproducibility as every earlier entry in this table). Section 26.10 records this as the same open item as before, narrowed once more rather than closed.

### 25.7.6 Does the curvature distribution flatten out, or converge to a fixed shape?

Section 25.7.2 established that curvature does not stay concentrated at the twelve pentagons as $n$ grows; the pentagon-corner vertices' share of the total defect fell from 76% at $n = 2$ to under 1% by $n = 20$. Left open there was what, if anything, the resulting distribution settles toward: whether refinement drives the sixty-vertex defect field toward the exactly uniform distribution a smooth round sphere's own curvature 2-form would suggest (every vertex eventually carrying an equal, vanishingly small share), or toward something else. This subsection computes the relevant statistics directly, at every frequency already tabulated plus $n = 50$, over *all* of $\mathcal{P}_n$'s vertices rather than the sixty pentagon-corner ones alone.

**Definition 25.14** (Three shape statistics)**.** For $\mathcal{P}_n$'s full set of Goldberg-vertex defects (plane-projected convention), let $\mathrm{CV}(n)$ be their coefficient of variation (population standard

deviation divided by mean), let $M(n)$ be the ratio of the maximum defect to the mean defect, and let $R(n)$ be the ratio of the sixty pentagon-corner vertices' mean defect to the mean defect of every other vertex.

**Proposition 25.15** (Computed at fourteen frequencies, $n = 2$ to 50)**.** *Table 25.5 reports* $\mathrm{CV}(n)$*,* $M(n)$*, and* $R(n)$ *at* $n = 2, 3, 4, 5, 6, 8, 10, 12, 16, 20, 30, 40, 50$.

| $n$ | 2 | 3 | 4 | 5 | 6 | 8 | 10 |
|---|---|---|---|---|---|---|---|
| $\mathrm{CV}(n)$ | 0.0190 | 0.0637 | 0.0881 | 0.1018 | 0.1102 | 0.1194 | 0.1241 |
| $M(n)$ | 1.011 | 1.077 | 1.131 | 1.157 | 1.145 | 1.186 | 1.194 |
| $R(n)$ | 1.0453 | 0.8861 | 0.8192 | 0.7777 | 0.7495 | 0.7136 | 0.6918 |
| $n$ | 12 | 16 | 20 | 30 | 40 | 50 | |
| $\mathrm{CV}(n)$ | 0.1267 | 0.1295 | 0.1309 | 0.1322 | 0.1327 | 0.1330 | |
| $M(n)$ | 1.190 | 1.201 | 1.203 | 1.204 | 1.206 | 1.206 | |
| $R(n)$ | 0.6772 | 0.6588 | 0.6477 | 0.6328 | 0.6253 | 0.6208 | |

Table 25.5: Definition 25.14's three shape statistics, all vertices of $\mathcal{P}_n$, plane-projected convention, at fourteen frequencies.

**Key Result**

*[The relative shape appears to converge to a fixed, nontrivial profile — not to uniformity] All three statistics move steadily across the tested range but toward what look like fixed nonzero, non-unit limits rather than toward the values —* $\mathrm{CV}(n) \to 0$ *and* $M(n) \to 1$ *— that exact uniformity would require: between* $n = 20$ *and* $n = 50$*, a two-and-a-half-fold increase in frequency,* $\mathrm{CV}(n)$ *moves only from* 0.1309 *to* 0.1330 *(a* 1.6% *change),* $M(n)$ *from* 1.203 *to* 1.206 *(a* 0.2% *change), and* $R(n)$ *from* 0.6477 *to* 0.6208 *(a* 4.2% *change, the largest of the three but still an order of magnitude slower than the pentagon-corner defect's own* 6.5*-fold drop over the same range). If the defect field were becoming uniform, all three statistics would keep moving toward* $\mathrm{CV} = 0$*,* $M = 1$*,* $R = 1$ *without bound; instead each appears to be approaching a plateau well short of those values. Taken together with section 25.7.5's* $n^{-2}$*-consistent overall decay, the data are consistent with a defect field of the form* $\delta_n(v) \approx \kappa(n) \cdot h(v)$*, where* $\kappa(n) \to 0$ *(at the rate section 25.7.5 narrows toward* $-2$*) is a single overall scale shared by every vertex and* $h$ *is a fixed, non-constant relative profile that refinement reveals rather than washes out.*

**Todo — verify primary source**

Section 25.7.6 is a reading of fourteen computed data points, not a proof that $\mathrm{CV}(n)$, $M(n)$, and $R(n)$ actually converge (as opposed to continuing to drift slowly, without bound, toward the uniform-limit values at a rate this chapter's range is too short to detect — exactly the same caveat section 25.7.5's todobox already gives for the defect magnitude itself). No candidate closed form for the limiting profile $h$ is proposed here; identifying one would require an analytic account of the specific discretization error introduced by this chapter's construction (barycentric subdivision of each icosahedral face, followed by centroid-taking and radial renormalization to the sphere, followed by definition 25.10's SVD-plane angle convention), which this chapter has not attempted. What this subsection changes is the shape of the open question itself: section 25.7's original framing — does the discrete singularity structure converge to a specific smooth vector field's zero set? — presupposes a limit supported on finitely many points, and section 25.7.6's evidence points the other way, toward every vertex retaining a comparable (relative) share of the shrinking total curvature rather than the curvature staying pinned to a fixed finite set of points as $n \to \infty$. If that reading holds up under further computation, the right limiting object is not a vector field's zero set at all but a continuous relative-density profile on $S^2$, and the singularity-convergence question section 26.10 has carried since chapter 25's first draft should be understood, from this point forward, in that revised and considerably sharper form.

### 25.7.7 An exact law for the mean defect, and what it explains

Sections 25.7.5 and 25.7.6 both read the $n^{-2}$ behavior off finite, computed data — a fit, not a derivation. This subsection derives one genuine $n^{-2}$ law outright, for a related but different quantity (the defect averaged over *every* vertex of $\mathcal{P}_n$, not the pentagon-corner value specifically), from two facts already established elsewhere in this volume, and then uses it to say something sharper about the pentagon-corner value after all.

**Proposition 25.16** (Vertex count, exactly)**.** *$\mathcal{P}_n$ has exactly $V(n) = 20n^2$ vertices, for every $n \geq 1$.*

*Proof.* Each of the icosahedron's twenty triangular faces is subdivided, in definition 25.5's class-I construction, by a barycentric grid of side $n$: the standard fact that such a grid partitions a triangle into $n^2$ congruent small triangles (one small triangle for each of the $n$ rows, contributing $1, 3, 5, \ldots, 2n-1$ upward- and downward-pointing triangles respectively, summing to $n^2$) is exactly what the verification script's `subdivide` routine implements, and is independent of which face is subdivided. Each small triangle's centroid is one vertex of $\mathcal{P}_n$ (definition 25.5), and distinct small triangles have distinct centroids (no two centroids coincide, since the construction's only identifications are along shared edges between adjacent icosahedral faces, which merge *vertices* of the fine triangulation, not centroids of its faces). Hence $V(n) = 20 \times n^2$. □

**Key Result**

*[The mean defect, exactly, for every $n$] Combining proposition 25.16 with proposition 25.7 (or, for the specific plane-projected convention used throughout section 25.7.2 onward, proposition 25.11), both of which give $\sum_v \delta(v) = 4\pi = 720^\circ$ exactly for every $n$:*

$$\bar{\delta}(n) \;:=\; \frac{1}{V(n)} \sum_{v} \delta(v) \;=\; \frac{720^\circ}{20n^2} \;=\; \frac{36^\circ}{n^2},$$

*exactly, for every $n \geq 1$ — not a fitted or extrapolated value, but an immediate corollary of two propositions already proved. At $n = 1$ (the dodecahedron) this reads $\bar{\delta}(1) = 36^\circ$, matching the dodecahedron's twenty identical $36^\circ$ vertex defects exactly, a useful sanity check on the formula at the one frequency where the mean and the individual value coincide by symmetry.*

**Key Result**

*[What this does and does not settle] Section 25.7.7 proves, rather than merely observing, that curvature decays exactly as* $n^{-2}$ on average*: the* $1/n^2$ *mechanism section 25.7.2 first floated as "a plausible mechanism, not a derived result" is now a derived result, for the mean. This does not by itself settle the pentagon-corner value's own asymptotics, since a mean can decay at a fixed rate while a particular vertex's value decays at a different rate, provided the two remain commensurate as* $n$ *grows — which is exactly section 25.7.6's open question, now sharpened by having an exact, rather than empirical, denominator to compare against.*

**Definition 25.17** (The exact ratio $\kappa(n)$)**.** Let $\kappa(n) := \delta_{\text{pent}}(n)/\bar{\delta}(n)$, where $\delta_{\text{pent}}(n)$ is tables 25.1 to 25.4's computed pentagon-corner defect (plane-projected convention) and $\bar{\delta}(n) = 36°/n^2$ is section 25.7.7's exact mean.

| $n$ | 2 | 3 | 4 | 5 | 6 | 8 | 10 |
|---|---|---|---|---|---|---|---|
| $\kappa(n)$ | 1.010963 | 0.921101 | 0.847983 | 0.799038 | 0.765461 | 0.723321 | 0.698273 |
| $n$ | 12 | 16 | 20 | 30 | 40 | 50 | |
| $\kappa(n)$ | 0.681769 | 0.661423 | 0.649391 | 0.633578 | 0.625777 | 0.621131 | |

Table 25.6: Definition 25.17's exact-denominator ratio $\kappa(n)$, at the same fourteen frequencies as table 25.5. Monotonically decreasing at every step tested, and — unlike table 25.5's $R(n)$, whose denominator is itself a computed statistic — computed against a closed-form, exact denominator.

**Proposition 25.18** ($\kappa(n)$ fits $\kappa_\infty + c/n$, sharply once the smallest frequencies are excluded)**.** *An ordinary least-squares fit of* $\kappa(n)$ *against* $1/n$*, using all thirteen of table 25.6's points (*$n = 2$ *to* 50*), gives* $\kappa_\infty = 0.6125$*,* $c = 0.860$ *(*$R^2 = 0.9863$*), with residuals as large as* 0.0317 *at* $n = 2$ *— a poor fit by this chapter's usual standard, and, as section 25.7.7 argues below, poor for an identifiable reason rather than merely by chance. Restricted to the seven points* $n \geq 10$*, the fit tightens to* $\kappa_\infty = 0.6015$*,* $c = 0.964$*, every residual under* $4 \times 10^{-4}$*; restricted further to the five points* $n \geq 16$*, it gives* $\kappa_\infty = 0.6021$*,* $c = 0.948$*, every residual under* $1.0 \times 10^{-4}$*; restricted to the three points* $n = 30, 40, 50$*, it gives* $\kappa_\infty = 0.6024$*,* $c = 0.934$*, every residual under* $1.5 \times 10^{-5}$*. Unlike propositions 25.12 and 25.13's log-log slope, whose point estimate keeps moving as the range narrows,* $\kappa_\infty$ *is already stable to three significant figures (*0.6015 *to* 0.6024*) across every range that excludes* $n < 10$ *— it is only the unrestricted fit, dragged toward* 0.6125 *by curvature at the smallest frequencies, that departs from this otherwise consistent value.*

**Key Result**

*[A single linear-in-$1/n$ correction accounts for most of the curvature in $\kappa(n)$ at large $n$, and, if $\kappa(n)$ converges, the pentagon-corner exponent is exactly $-2$] The full-range fit's own residual pattern is informative:* $0.0317$ *at* $n = 2$*, shrinking and changing sign by* $n = 10$*, then settling to residuals of order* $1 \times 10^{-4}$ *or smaller for* $n \geq 16$ *(proposition 25.18) is exactly the signature of a subleading term — a* $1/n^2$ *or higher correction on top of the fitted* $\kappa_\infty + c/n$ *— that matters at small* $n$ *and becomes negligible as* $n$ *grows, rather than evidence against the linear-in-*$1/n$ *model itself. Once that small-*$n$ *curvature is excluded, the three restricted fits' residuals — at most a few times* $1 \times 10^{-4}$ *over* $n \geq 10$*, tightening to a few times* $1 \times 10^{-5}$ *over* $n = 30, 40, 50$ *— are far smaller than section 25.7.4's own U-shaped residuals from fitting* $\log(\text{defect})$ *directly against* $\log n$*, which is exactly what a hidden* $1/n$ *correction term predicts: if* $\delta_{\text{pent}}(n) = \bar{\delta}(n)\,\kappa(n) \approx \frac{36^\circ}{n^2}\left(\kappa_\infty + \frac{c}{n}\right) = \frac{36^\circ \kappa_\infty}{n^2} + \frac{36^\circ c}{n^3} + \cdots$*, then* $\log \delta_{\text{pent}}(n)$ *is not exactly linear in* $\log n$ *at any finite* $n$ *— it has slope* $-2$ *only in the strict* $n \to \infty$ *limit, approaching it from below at a rate set by* $c/(\kappa_\infty n)$*, which is precisely the slow, monotone approach sections 25.7.3 and 25.7.5 already found empirically without identifying its source. Taking the large-*$n$ *restricted fits' extrapolation at face value —* $\kappa(n) \to \kappa_\infty \approx 0.602$*, a genuine finite nonzero limit, rather than* $\kappa(n) \to 0$ *or* $\kappa(n) \to \infty$ *— gives* $\delta_{\text{pent}}(n) \sim 36^\circ \kappa_\infty\, n^{-2} \approx 21.7^\circ \cdot n^{-2}$ *as* $n \to \infty$*: the exponent is exactly* $-2$*, not a nearby value, with section 25.7.5's* $-2.049$*-and-still-moving log-log slopes fully explained as the finite-*$n$ *signature of the same* $1/n$ *correction the restricted fits capture directly. The unrestricted fit's poor residuals do not undermine this reading; they are consistent with it, once read as evidence that a single linear term is not the whole story at small* $n$*, only at large* $n$*.*

**Todo — verify primary source**

Section 25.7.7's conclusion — exponent exactly $-2$ — rests on the large-$n$ restricted fits' extrapolated intercept, not on a proof that $\kappa(n)$ converges at all, nor on a derivation of $\kappa_\infty$'s value from the construction's geometry (why 0.602, specifically, rather than some other constant between 0 and 1, is not explained here; it is read off a fit, exactly as proposition 25.12's slope was). Nor does the unrestricted, full-range fit's poor quality — residuals two orders of magnitude larger than the restricted fits, and an extrapolated $\kappa_\infty = 0.6125$ visibly displaced from the restricted fits' 0.602 — follow from anything derived here: it is read directly off table 25.6, and its cause (curvature beyond the fitted $1/n$ term, concentrated at the smallest frequencies) is a plausible diagnosis, not a proven decomposition of $\kappa(n)$ into a linear term plus a specific higher-order correction. What has changed relative to section 25.7.5's todobox is the *quality* of the evidence at large $n$: a three-parameter question (does $\kappa(n)$ converge, and if so to what, with what correction), restricted to $n \geq 16$ or further to $n = 30, 40, 50$, fits the data to residuals near the level of numerical noise, where proposition 25.12's one-parameter question (what is the log-log slope) did not even at its most restricted range, and section 25.7.7 independently guarantees the $n^{-2}$ scaling for the *mean*, removing one degree of freedom from the fit rather than leaving it purely empirical. This is markedly stronger evidence for exponent $-2$ than anything sections 25.7.4 and 25.7.5 produced, short of an actual proof that $\lim_{n\to\infty} \kappa(n)$ exists — but the evidence is genuinely stronger only at large $n$; a naive reader of table 25.6 alone, fitting the full range without restriction, would not reach the same confidence. Section 26.10 should record the asymptotic-rate item in this strengthened form: not "consistent with $-2$, confidence interval not yet reaching it," but "consistent with $-2$ via an exact $n^{-2}$ law for the mean and a ratio $\kappa(n)$ that extrapolates cleanly to $\approx 0.602$ once frequencies below $n = 10$ are excluded, with the existence of that limit, and an account of why the small-$n$ fit misbehaves, the two remaining gaps."

### 25.7.8 Is $\kappa_\infty$ scheme-invariant? A three-way construction test

Before searching further for a closed form, a prior question deserves an answer: is $\kappa_\infty \approx 0.60282$ a property of the icosahedral five-fold disclination itself, or an artifact of this chapter's particular mesh-generation scheme — single-shot barycentric interpolation of each icosahedron face, $(i \cdot A + j \cdot B + k \cdot C)/n$, followed by one radial projection onto the unit sphere ("SSB" below)? `scripts/kappa_scheme_comparison.py` tests this directly, by implementing two

further, genuinely different class-I subdivision schemes and repeating the same pentagon-corner-defect-under-plane-projected-convention computation and quadratic-in-$1/n$ extrapolation for each:

- **RMS** (recursive midpoint bisection): at every subdivision *level*, each new edge midpoint is renormalized to the unit sphere before the next level is built, reaching only $n = 2^k$. Since $\mathrm{slerp}(p, q, \frac{1}{2}) = (p + q)/\|p + q\|$ for unit vectors $p, q$, this scheme already uses exact geodesic midpoints at every level, unlike SSB's single-shot chords.

- **GEO** (recursive two-step slerp): a spherical analogue of a barycentric triangle, reducing to true great-circle points along every boundary edge and at every interior grid point, for arbitrary $n$.

**Epistemic Notice**

Before trusting any of the three extrapolated values, the exact Gauss–Bonnet identity (total plane-projected defect over all Goldberg vertices = 720°) was checked for every mesh in all three schemes at every tested frequency; it holds to $1 \times 10^{-6}$ degrees or better in every case, confirming that all three constructions are genuinely valid flat-faced polyhedra combinatorially equivalent to a sphere, not merely superficially similar meshes.

**Key Result**

*[$\kappa_\infty$ is not scheme-invariant: three constructions, three different limits] Quadratic-in-$1/n$ extrapolation (the same procedure as proposition 25.36, over comparable ranges) gives:*

$\kappa_\infty^{\mathrm{SSB}} \approx 0.602711$ ($n$=10–50), $\kappa_\infty^{\mathrm{RMS}} \approx 0.923659$ ($n$=8–64), $\kappa_\infty^{\mathrm{GEO}} \approx 0.951982$ ($n$=10–50),

*all three fits with maximum residual under* $1.3 \times 10^{-5}$ *— comparably clean convergence in every case, not an artifact of a poorly-behaved fit in one of the three. SSB's value reproduces section 25.7.28's independently, far more extensively cross-checked* 0.60282 *to four significant figures, confirming this reimplementation is faithful to the construction used throughout the rest of this chapter. RMS and GEO, however, converge cleanly to values roughly* 50–60% *larger than SSB's, and to each other only within about* 3% *— not a near-miss consistent with numerical noise, but three genuinely different limits.*

**Epistemic Notice**

The honest reading of section 25.7.8 is negative but informative: $\kappa_\infty$ is not an invariant of the icosahedral five-fold vertex alone — it is a construction-specific quantity, sensitive to exactly how curvature is distributed between the mesh-generation step and the final radial projection. This sharpens, with a concrete counterexample rather than only an a-priori argument, section 25.7.20's earlier reading ("$\kappa_\infty$ measures a property of *this specific* construction's discretization error …there is no a priori reason a quantity defined that way should reduce to a simple algebraic number at all"): it also forecloses the hope that a scheme-independent closed form might exist and simply not yet be found by proposition 25.37's candidate screen. Any future closed-form derivation — via section 25.7.21's Feinberg–Friedberg–Lee–Ren route or otherwise — would necessarily be a closed form for *this particular* SSB construction's discretization error, not for the disclination geometrically. RMS and GEO's mutual closeness (within 3%) relative to their shared distance from SSB (∼55%) is suggestive — both are genuinely geodesic constructions (exact great-circle interpolation, at every level for RMS, at every grid point for GEO), while SSB alone uses flat Euclidean chords before a single final projection — but this chapter does not attempt to explain the RMS/GEO gap itself, nor to identify which geometric feature of SSB specifically produces its markedly different limit; both are new open items, narrower than the pre-existing closed-form question, that a future pass could take up.

### 25.7.9 Isolating the SSB/RMS/GEO gap: flat chords or nested normalization?

Section 25.7.8's own closing disclaimerbox left the RMS/GEO gap from SSB unexplained, naming two features that differ simultaneously between SSB on one side and RMS/GEO on the other: SSB uses flat Euclidean chords between the original face corners, computed and normalized once per grid point, while RMS and GEO both use geodesic interpolation (exact for RMS's midpoints, `slerp` for GEO) applied repeatedly — at every bisection level for RMS, at each of two nested calls per grid point for GEO. `scripts/kappa_hybrid_scheme.py` isolates these two axes directly, by constructing a fourth scheme, HYB, that reuses GEO's exact two-step recursive grid construction (interpolate first along edge $BC$, then from $A$ to that intermediate point) but replaces each of GEO's two `slerp` calls with a flat linear interpolation followed by a renormalization: flat chords, like SSB, but nested, twice-per-point normalization, like GEO, rather than SSB's single normalization computed directly from $A, B, C$.

**Epistemic Notice**

As before, the exact Gauss–Bonnet identity (total plane-projected defect = 720°) was checked for HYB at every tested frequency; it holds to $1 \times 10^{-6}$ degrees, confirming HYB is as valid a flat-faced polyhedral construction as the other three.

**Key Result**

*[The SSB/RMS/GEO gap tracks flat-vs-geodesic interpolation, not single-shot-vs-nested normalization] Quadratic-in-1/n extrapolation ($n = 10$–$50$, the same range and procedure as SSB and GEO above) gives*

$$\kappa_\infty^{\mathrm{HYB}} \approx 0.626743 \quad (\textit{max residual } 9.6 \times 10^{-6}),$$

*about 4% above $\kappa_\infty^{\mathrm{SSB}} \approx 0.602711$, and roughly 47% and 52% below $\kappa_\infty^{\mathrm{RMS}} \approx 0.923659$ and $\kappa_\infty^{\mathrm{GEO}} \approx 0.951982$ respectively (relative to HYB itself). HYB shares SSB's flat-chord interpolation geometry and GEO's nested, twice-per-point renormalization structure; its limit sits with the scheme it shares interpolation geometry with, not the scheme it shares normalization structure with — and the SSB/HYB gap (∼4%) is comparable in size to RMS and GEO's own ∼3% mutual gap, both far smaller than the ∼50–60% that separates either pair from the other.*

**Epistemic Notice**

The honest reading: the large gap section 25.7.8 found between SSB and RMS/GEO is driven overwhelmingly by whether the underlying interpolation is a flat Euclidean chord or a true spherical geodesic, not by how many times normalization is applied during construction — removing SSB's single-shot normalization while keeping flat chords (HYB) barely moves $\kappa_\infty$. This narrows, rather than resolves, the open item section 25.7.8's disclaimerbox left: the mechanism responsible for the gap is now identified (flat interpolation systematically under- or over-curves relative to true geodesics, in a way normalization frequency cannot compensate for), but a quantitative account of *why* flat interpolation produces a smaller limiting defect specifically at a five-fold vertex, as opposed to a plausible qualitative story, is not attempted here and remains open — along with the pre-existing, now-narrower question of why RMS and GEO, both geodesic, still differ from each other by ∼3% despite sharing the same interpolation geometry.

### 25.7.10 Why RMS and GEO still differ: nested slerp does not commute

Section 25.7.9's own closing disclaimerbox left one item narrower but still open: RMS and GEO, sharing the same flat-vs-geodesic axis (both fully geodesic) and differing only in how their geodesic interpolation is composed — RMS by repeated bisection across levels, GEO by a single two-step `slerp` evaluated directly from each grid point's combinatorial fraction —

nonetheless converge to visibly different values, $\kappa_\infty^{\mathrm{RMS}} \approx 0.923659$ against $\kappa_\infty^{\mathrm{GEO}} \approx 0.951982$, a $\sim 3\%$ mutual gap. `scripts/rms_geo_noncommute.py` tests a specific candidate mechanism directly, rather than only naming it qualitatively: that RMS and GEO are not two algorithms for computing the same grid, but two grids that are genuinely different as maps from barycentric fraction to the sphere, because `slerp` does not commute with itself under repeated application the way flat affine interpolation does. In the flat-triangle case, bisecting each edge $k$ times and reading off the point combinatorially labelled $(i, j, k)$ with $i + j + k = 2^{\mathrm{level}}$ is identical to the direct barycentric point $(iA + jB + kC)/n$ — affine interpolation is associative, so nested bisection and direct barycentric averaging agree exactly, for every level and every label. On the sphere there is no a-priori reason this identity should survive, since `slerp` is not affine.

**Proposition 25.19** (RMS and GEO disagree pointwise at every fixed barycentric fraction, and the disagreement does not vanish under refinement)**.** *On one actual icosahedron face* $(A, B, C)$*, for a fixed rational fraction* $(p, q, r)$ *with* $p + q + r = 1$ *realized as* $(i, j, k) = (pn, qn, rn)$ *at every* $n = 2^{\mathrm{level}}$ *for which this is integral, the chord distance between RMS's point (built by section 25.7.9's recursive-bisection construction, in which a point created at one level is never revisited at finer levels) and GEO's point (built by section 25.7.9's direct two-step* `slerp`*, evaluated from the same fraction* $(p, q, r)$ *regardless of* $n$*) is* constant in $n$ *at that fraction, not merely small:* $0.044\,959\,59$ *for the specific fraction* $(1/4, 1/4, 1/2)$*, checked at* $n = 4, 8, 16, 32, 64, 128$ *and agreeing to eight decimal places at every one. Averaged over all grid labels present at each level, the mean RMS/GEO discrepancy rises from* $0.005\,33$ *at* $n = 4$ *to* $0.013\,82$ *at* $n = 128$*, still increasing but visibly flattening (successive increments* $0.0040, 0.0023, 0.0013, 0.0007, 0.0003$*), consistent with convergence to a nonzero asymptotic value rather than to zero.*

**Epistemic Notice**

The exact constancy of the single-fraction discrepancy across six frequencies is not a coincidence to be explained further: it follows immediately from how each construction is defined. GEO's point at fraction $(p, q, r)$ is, by section 25.7.9's own formula, a function of $(p, q, r)$ alone (the two `slerp` calls use only the ratios $k/(j{+}k)$ and $(j{+}k)/n$, both scale-invariant), so it is literally the same point at every $n$ for which that fraction is realized. RMS's point at the same fraction is fixed once the bisection level that first produces it is reached, and *never recomputed* at finer levels, by construction (finer bisection only adds new points between existing ones, never revisits old vertices). The two points being compared at $n = 4, 8, \ldots, 128$ are therefore, after the first level that realizes the fraction, literally the same pair of points every time — what proposition 25.19 establishes is not that this particular repeated comparison is interesting in itself, but that the pair is nonzero at all: RMS and GEO's respective limiting maps from barycentric fraction to $S^2$ genuinely disagree, pointwise, rather than merely differing by a refinement-order discretization error that both would eventually converge past.

**Key Result**

*[A mechanism for the RMS/GEO gap, not yet a quantitative derivation of it] Proposition 25.19 establishes something section 25.7.9's closing disclaimerbox could not: RMS and GEO do not converge to the same continuous map of the spherical triangle as $n \to \infty$. Both are legitimate, fully geodesic constructions in the local sense (every edge in both is a true great-circle arc), but they parametrize the triangle's interior differently, and that difference is exactly the kind of thing that would shift a defect computed from local vertex neighborhoods — as $\kappa_\infty$ is — by a fixed relative amount rather than driving it to a common limit. This is a candidate mechanism, identified and directly demonstrated rather than merely conjectured, for the $\sim 3\%$ RMS/GEO gap section 25.7.9 left unexplained. It is not yet a quantitative derivation of that specific 3%: proposition 25.19's measured discrepancies (chord distances of order $1 \times 10^{-2}$ to $4 \times 10^{-2}$ on a unit sphere, at a single representative face and a handful of fractions) have not been propagated through the pentagon-corner-defect calculation itself to check that they produce a gap of the right sign and magnitude, only shown to be real, nonvanishing, and of a plausible order. Doing so — relating the local Jacobian of RMS's versus GEO's parametrization at the pentagon corners specifically to the $\sim 3\%$ defect-ratio gap — remains open, and is a more tractable, better-specified version of the question section 25.7.9 left than the original "why do RMS and GEO differ" was.*

### 25.7.11 Propagating the mechanism: why GEO's $\kappa_\infty$ is not one local number

Section 25.7.10's closing text left one item open: relating the local Jacobian of RMS's versus GEO's parametrization at the pentagon corners specifically to the $\sim 3\%$ defect-ratio gap. This subsection closes it, computationally, and finds an answer more specific than the question anticipated.

**Proposition 25.20** ($\delta(v)$ at a pentagon corner is a function of exactly four points, and only one of them is scheme-dependent)**.** *Let $v$ be the Goldberg vertex at a pentagon corner, i.e. the centroid of the corner triangle $(A, P, Q)$ where $A$ is an original icosahedron vertex and $P, Q$ are the grid points at labels $(n-1, 1, 0)$, $(n-1, 0, 1)$. Under the raw-angle convention (section 25.7.1), $\delta(v)$ depends only on $v$ and its three neighbour centroids $N_{\text{left}}$, $N_{\text{right}}$, $N_{\text{int}}$:*

$$\delta(v) = 360^\circ - \angle(N_{\text{left}}, v, N_{\text{right}}) - \angle(N_{\text{right}}, v, N_{\text{int}}) - \angle(N_{\text{int}}, v, N_{\text{left}}).$$

*$N_{\text{left}}$ and $N_{\text{right}}$ (the corner-triangle centroids of the two icosahedral faces adjacent across edges $AP$, $AQ$) equal $v$ rotated by $\mp 72^\circ$ about the axis through $A$, exactly, for any scheme built from geodesics – verified computationally to floating-point precision ($< 2 \times 10^{-16}$) against an explicit two-face mesh. $N_{\text{int}}$, the centroid of the one remaining adjacent triangle $(P, Q, R)$ with $R$ at label $(n-2, 1, 1)$, is the only one of the four points that can differ between RMS and GEO, since $P, Q$ are themselves pure single-edge slerp values (identical for both) and $R$ is the first grid point off the boundary.*

*Verification.* `scripts/pentagon_corner_local_jacobian.py` computes $\delta(v)$ two ways at matching $n$ – once via this four-point reduction, once via an explicit full mesh grid dict – and finds exact agreement (to floating-point precision) at every $n$ tested. The rotation identity for $N_{\text{left}}, N_{\text{right}}$ is checked directly against a two-face patch of the icosahedron built from scratch, independent of scheme. □

**Epistemic Notice**

A separate script, `scripts/pentagon_corner_jacobian_fullmesh_crosscheck.py`, repeats this check a second, independent way: rather than reusing any grid-construction code from `pentagon_corner_local_jacobian.py` itself, it builds a complete GEO mesh of the whole icosahedron with `subdivide_geo()`, imported unchanged from section 25.7.8's `kappa_scheme_comparison.py` (written for the unrelated SSB/RMS/GEO scheme-invariance question, and never previously used for this purpose), extracts the role-A pentagon corner's full vertex star by shared-edge adjacency alone, and computes $\delta(v)$ from that complete, honestly-assembled star under the same raw-angle convention. At $n = 4, 6, 8$ this matches the four-point reduction's value to floating-point precision ($\leq 5.7 \times 10^{-14}$ absolute difference in every case), confirming proposition 25.20 against a mesh built by genuinely different code, not merely against `pentagon_corner_local_jacobian.py`'s own internal grid dict.

**Key Result**

*[The four-point reduction reproduces RMS's $\kappa_\infty$, but not GEO's – because GEO is not one local geometry] Applying proposition 25.20 with RMS's $R$ and extrapolating $\kappa(n)$ the same way as proposition 25.18 gives $\kappa_\infty^{\mathrm{RMS}} \approx 0.9277$, within 0.4% of the chapter's independently, globally computed 0.923659 — confirming the reduction. Doing the same with GEO's $R$ gives $\kappa_\infty \approx 1.0135$: not close to 0.951982, and on the wrong side of it. The reason is structural, not numerical: RMS's midpoint bisection is manifestly symmetric under permuting a face's three vertices, so every pentagon corner sees the same local geometry. GEO's formula, $\mathrm{geo}(i, j, k) = \mathrm{slerp}(A, \mathrm{slerp}(B, C, k/(j+k)), (j+k)/n)$, is not symmetric in $(A, B, C)$: it privileges $A$ as the outer `slerp`'s start point. A single face therefore has three pentagon-corner geometries, not one – one at the privileged corner ("role A") and two, identical to each other by the formula's own $B \leftrightarrow C$ symmetry, at the other two ("role B/C"). Computed separately (same extrapolation protocol, $n = 64$ to 1024):*

$$\kappa_\infty^{role\ A} \approx 1.013507, \qquad \kappa_\infty^{role\ B=C} \approx 0.927697.$$

*The $(1, 2)$-weighted average across a face's three corners, $(\kappa^A + 2\kappa^B)/3 \approx 0.9563$, reproduces the globally computed $\kappa_\infty^{\mathrm{GEO}} = 0.951982$ to within about 0.4% – the same order as the raw-angle/plane-projected convention gap section 25.7.2 already found negligible, and consistent with it being the residual source of the small remaining discrepancy. GEO's reported $\kappa_\infty$ is therefore not a single scheme-vs-scheme constant to be compared against RMS's; it is a mixture of two distinct local behaviours, one of which ($\kappa^{B=C} \approx 0.9277$) is itself already close to RMS's own value.*

**Key Result**

*[The local Jacobian, isolated: sign and order of magnitude confirmed] Since $P, Q$ are shared between every scheme and role, the entire local Jacobian difference reduces to where each construction places the single point $R$. Measuring $R$'s chord displacement from the flat bisector $M = \mathrm{midpoint}(P, Q)$ (the only quantity a shared linear/tangent-plane approximation cannot see, since $P, Q$ fix that part identically for every case) gives, at matching $n$, a stable ratio*

$$\frac{|R_{role\ A} - M|}{|R_{role\ B} - M|} \longrightarrow 1.2713 \quad (n = 256 \text{ to } 16384, \text{ stable to 4 significant figures}),$$

*against $\kappa_\infty^{role\ A}/\kappa_\infty^{role\ B} = 1.0925$. Both displacements point in the same direction (outward from $A$, away from the boundary chord, for every $n$ tested) at every role, so the mechanism predicts the correct sign of the gap (role A's construction curves further from the flat bisector than role B's, and role A does carry the larger $\kappa_\infty$) and the correct order of magnitude (a $\sim 27\%$ larger local displacement producing a $\sim 9\%$ larger $\kappa$, both $O(10\%)$ effects, not $O(1\%)$ or $O(100\%)$ ones) – but the two ratios are not equal, confirming that $\delta(v)$ is a genuinely geometric (shape-dependent) function of $R$'s position, not a linear functional of its displacement magnitude alone.*

**Todo — verify primary source**

Section 25.7.11 answers section 25.7.10's open item at the level it was posed – sign and order of magnitude, propagated through the actual pentagon-corner-defect calculation rather than only demonstrated as a nonvanishing chord discrepancy – but does not derive a closed-form map from $R$'s displacement to $\kappa_\infty$'s numerical value; the $1.2713 \to 1.0925$ relationship is read off computed data, not derived from the angle-defect formula analytically. It also leaves a new, narrower question in its place: section 25.7.11's finding that GEO's $\kappa_\infty$ is a $(1, 2)$-weighted mixture of two role-specific constants is a property of *this particular* two-step `slerp` formula's asymmetry, not of geodesic parametrization in general – RMS carries no such mixture because its construction is manifestly permutation-symmetric. Whether a genuinely $(A, B, C)$-symmetric geodesic two-step construction exists that would collapse GEO's three roles into RMS's one, and if so whether it would then match RMS's $\kappa_\infty$ or a third value, is not addressed here.

### 25.7.12 A genuinely symmetric construction, built and tested: it does not land on RMS

The previous todobox's question is taken up directly here, by constructing the symmetric alternative it asked for rather than continuing to speculate about it.

**Definition 25.21** (SYM: the spherical Fréchet-mean grid)**.** For weights $(i/n, j/n, k/n)$ summing to 1 assigned to $(A, B, C)$, define the SYM grid point as the spherical (Karcher/Fréchet) weighted mean of Buss and Fillmore (2001, "Spherical averages and applications to spherical splines and interpolation"): the point $x \in S^2$ minimizing the weighted sum of squared geodesic distances to $A, B, C$, computed by iterating $x \leftarrow \mathrm{Exp}_x\big(\sum_v w_v \,\mathrm{Log}_x(v)\big)$ from the normalized Euclidean weighted average. Unlike GEO's $\mathrm{slerp}(A, \mathrm{slerp}(B, C, \cdot), \cdot)$, this expression treats $(A, w_A)$, $(B, w_B)$, $(C, w_C)$ identically: permuting the three (vertex,weight) pairs simultaneously leaves $x$ unchanged by construction, not merely by an unproved symmetry of the numbers it happens to produce.

**Proposition 25.22** (SYM reduces to ordinary slerp on each edge, and is exactly permutation-symmetric)**.** *On an edge ($k = 0$), SYM's weighted mean of two points reduces to the ordinary two-point slerp at the corresponding fraction: checked directly against definition 25.21's iteration at four fractions along edge $AB$ ($n = 16$), agreeing with* $\mathrm{slerp}(A, B, \cdot)$ *to $2.6 \times 10^{-16}$ or better in every case. Applying proposition 25.20's pentagon-corner defect formula to SYM and anchoring it at each of $A$, $B$, $C$ in turn (i.e. computing the same corner via section 25.7.11's formula, with the roles of the three vertices cyclically permuted) gives $\delta(v) = 0.0328223136°$ at $n = 32$ in all three cases, with a reported maximum pairwise difference of exactly $0°$: SYM is not merely close to permutation-symmetric, as GEO's own reported $B = C$ agreement already suggested for two of its three roles, but exact to floating-point precision across all three.*

*Verification.* `scripts/kappa_symmetric_scheme.py` implements definition 25.21's iteration directly (50 Karcher-mean steps, convergence tolerance $1 \times 10^{-15}$ on the tangent-vector norm) and checks both claims computationally: the edge-reduction sanity check, and the three-anchor permutation-symmetry check, exactly as described above. □

**Key Result**

*[SYM's $\kappa_\infty$ coincides with GEO's non-privileged role, not with RMS] Extrapolating $\kappa(n)$ for SYM the same way as every other scheme in this chapter, at $n = 64$ to 8192 (eight frequencies, one full order of magnitude further than section 25.7.11's $n = 64$-to-1024 range), gives a value stable to five significant figures across every model order and range tested: linear extrapolation over all eight points gives 0.927703; quadratic and cubic extrapolation, over all eight points and over every truncated range tried down to $n \geq 512$, agree with each other to six decimal places at $\kappa_\infty^{\mathrm{SYM}} = 0.927702$. Compared against the two candidates section 25.7.11's closing todobox named: the gap to GEO's non-privileged role, $\kappa_\infty^{role\ B=C} = 0.927697$ (section 25.7.11), is $4.77 \times 10^{-6}$ — within the noise of the extrapolation fit itself, and three orders of magnitude smaller than the gap to RMS's own $\kappa_\infty^{\mathrm{RMS}} = 0.923659$, which is $4.04 \times 10^{-3}$, roughly 0.4% relative and far outside that noise floor. SYM, the construction built specifically to be exactly symmetric under permuting $A, B, C$, does not collapse onto RMS — itself also exactly permutation-symmetric, by a different mechanism (midpoint bisection rather than a weighted Fréchet mean) — but instead lands, to within extrapolation noise, on the value GEO's own asymmetric two-step formula already produces at its two non-privileged corners.*

**Todo — verify primary source**

Section 25.7.12 answers section 25.7.11's closing question at the level it was posed — a genuinely $(A, B, C)$-symmetric geodesic two-step construction does exist, and it does not reduce the three-role picture to RMS's value — but the specific answer it returns is not the "third value" the question anticipated either: $\kappa_\infty^{\mathrm{SYM}}$ agrees with $\kappa_\infty^{\mathrm{role\ B=C}}$ to a precision ($4.77 \times 10^{-6}$) that looks more like coincidence-worth-explaining than like two independently unrelated numbers landing near each other by chance, especially set against the $4.04 \times 10^{-3}$ gap to RMS that a naive "all symmetric constructions should agree" expectation would have predicted to be the small one instead. No mechanism for this near-coincidence is proposed here: definition 25.21's global weighted Fréchet mean and GEO's two-step nested `slerp`, evaluated at a non-privileged corner, are different functions of $(A, B, C, i, j, k)$ by construction, and nothing in proposition 25.22's verification explains why their respective pentagon-corner limits should agree to five significant figures while RMS's — built, like SYM, to be exactly symmetric — does not. Whether the agreement is exact (a genuine identity between the two constructions' local geometry at a non-privileged corner, perhaps provable directly from definition 25.21's Karcher-mean fixed-point equation) or merely close is left open, as is, unchanged from before, any explanation of the residual $\sim 3\%$ RMS/GEO gap itself.

### 25.7.13 Sharpening the coincidence: it already holds at the level of the raw displacement, not just the extrapolated $\kappa_\infty$

Section 25.7.12's finding was stated at the level of the extrapolated limit $\kappa_\infty$ alone. Section 25.7.11 already established, for GEO's own two roles, that a $\kappa_\infty$ comparison and a raw chord-displacement comparison need not agree (the 1.2713 displacement ratio against the 1.0925 $\kappa$ ratio): matching $\kappa_\infty$ values do not, by themselves, imply matching local geometry. This subsection checks the stronger, more specific claim directly, at the level of $R$'s own position rather than only the defect it produces.

**Proposition 25.23** (SYM's displacement magnitude converges to GEO's role-B value, not role A's, at rate $O(1/n)$)**.** *Using proposition 25.20's notation, let $d_{\mathrm{SYM}}(n) = |R_{\mathrm{SYM}} - M|$ be SYM's chord displacement at the $A$-anchored corner (label $(n-2, 1, 1)$, the same label GEO's own role-A $R$ uses), and let $d_{role\ A}(n)$, $d_{role\ B}(n)$ be GEO's own two role displacements as in section 25.7.11. Computed at $n = 256, 1024, 4096, 16384, 65536$:*

$$\frac{d_{\mathrm{SYM}}(n)}{d_{role\ A}(n)} \to 0.78658, \qquad \frac{d_{\mathrm{SYM}}(n)}{d_{role\ B}(n)} \to 1$$

*with the second ratio's deviation from* 1 *shrinking by very close to a factor of* 4 *at every quadrupling of* $n$ *tested (*$1.195 \times 10^{-3}, 3.01 \times 10^{-4}, 7.5 \times 10^{-5}, 1.9 \times 10^{-5}, 5 \times 10^{-6}$ *at* $n = 256, 1024, 4096, 16384, 65536$ *respectively) — clean* $O(1/n)$ *convergence to an exact match, not merely a stable small residual as the first ratio shows. SYM's* $R$*, though built from the same* $(n{-}2, 1, 1)$ *label GEO assigns to its own* privileged *corner, converges in raw magnitude to GEO's* non-*privileged corner's displacement, not to GEO's own same-label value.*

*Verification.* A new diagnostic script computes, at each tested $n$, both role-A's shared points $(P, Q, M, v$ at the $A$-anchored corner, identical for every scheme by proposition 25.20) and role-B's own $(P_B, Q_B, M_B, v_B$ at the $B$-anchored corner), places $R_{\text{SYM}}$ via definition 25.21's iteration at label $(n{-}2, 1, 1)$, $R_{\text{role A}}$ and $R_{\text{role B}}$ via GEO's own formula at their respective labels $(n{-}2, 1, 1)$ and $(1, n{-}2, 1)$, and reports the three chord distances and their ratios directly, exactly as tabulated above. A second check confirms the displacement is essentially radial in every one of the three cases at every $n$ tested: the unit vector $(R - M)$ has dot product $-1.000000$ (to six decimals already at $n = 256$) with the unit vector $(v - M)$, i.e. $R$, $M$, $v$ are collinear to the precision reported, with $R$ on the far side of $M$ from $v$ (outward, consistent with section 25.7.11's qualitative reading) — so the magnitude match found here is not an artifact of comparing displacements pointing in different directions. □

**Key Result**

*[The SYM/GEO-role-B agreement is a local, not merely an extrapolated, phenomenon] Proposition 25.23 sharpens section 25.7.12's finding in a specific direction: the coincidence is not an artifact of extrapolation smoothing away finite-n differences that happen to average out to the same limit — it already holds, to a shrinking residual consistent with an exact* $n \to \infty$ *identity, at the level of a single grid point's raw displacement at every finite* $n$ *tested from* 256 *up. This rules out one class of explanation section 25.7.12's closing todobox left open (that the* $\kappa_\infty$ *match might be a coincidence specific to how the extrapolation happens to weight finite-n data) without supplying a replacement: no derivation is given here for* why $\text{Exp}_x(\sum_v w_v \text{Log}_x(v))$*'s fixed point at weights* $((n{-}2)/n, 1/n, 1/n)$ *should converge, in raw displacement from the shared midpoint* $M$*, to the same value as* $\text{slerp}(A, \text{slerp}(B, C, \cdot), \cdot)$ *evaluated at the differently-labeled, differently-anchored weights* $(1/n, (n{-}2)/n, 1/n)$*. The two constructions are being compared at genuinely different barycentric labels here (SYM's own dominant vertex is A; GEO's role-B dominant vertex is B), related only by the ambient icosahedral symmetry that makes the two corners' local geometries comparable at all.*

**Todo — verify primary source**

What remains open is unchanged in kind from section 25.7.12's own closing todobox, only narrowed to a more specific and more clearly well-posed statement: a direct algebraic argument, from definition 25.21's Karcher-mean fixed-point equation and GEO's explicit two-`slerp` formula, for why the two constructions' respective $R$-placements coincide asymptotically at these particular, differently-anchored labels — or, short of a full derivation, at least a leading-order ($O(1/n)$) perturbative expansion of each construction around the sphere's dominant-vertex limit, matched term by term, which would explain the observed $O(1/n)$ convergence rate directly rather than only exhibiting it numerically. Neither is attempted here. The stable but unexplained constant $d_{\text{SYM}}/d_{\text{role A}} \to 0.78658$ is recorded for completeness and is not screened against any closed-form candidate list in this pass.

### 25.7.14 The mechanism isolated directly: privilege/nearness mismatch, not vertex identity, drives the split

Section 25.7.13's closing todobox left the comparison resting on icosahedral symmetry: SYM's displacement is checked at the $A$-anchored label $(n{-}2, 1, 1)$, GEO's role-B displacement at the differently-anchored label $(1, n{-}2, 1)$, and the two are only comparable at all

because the ambient symmetry carries one corner to the other. This subsection removes that step. Rather than comparing two constructions at two different corners, a single family of constructions is evaluated at the *same* label, varying only which vertex the formula itself treats as privileged.

**Definition 25.24** (Base-swapped GEO variant)**.** For a choice of *base vertex* $\beta \in \{A, B, C\}$, define $\mathrm{geo}_\beta(i, j, k) := \mathrm{slerp}\big(\beta,\ \mathrm{slerp}(\gamma_1, \gamma_2,\ w_2/(w_1{+}w_2)),\ (w_1{+}w_2)/n\big)$, where $\{\gamma_1, \gamma_2\} = \{A, B, C\} \setminus \{\beta\}$ and $w_1, w_2$ are the barycentric weights of $\gamma_1, \gamma_2$ respectively (so $\mathrm{geo}_A$ is exactly GEO's own official formula as used throughout this chapter, with $A$ always in the privileged slot). Every $\mathrm{geo}_\beta$ is built from the identical two-`slerp` recipe; only the label of which vertex sits in the privileged, outer-`slerp` position changes.

**Proposition 25.25** (At a fixed corner, only the base label matters, and it matters in exactly two ways)**.** *Fix the $A$-anchored pentagon corner: $P, Q, \mathrm{apex}$ at labels $(n{-}1, 1, 0), (n{-}1, 0, 1), (n, 0, 0)$, and $R$ at label $(n{-}2, 1, 1)$, all evaluated under $\mathrm{geo}_\beta$ for a chosen base $\beta$. By proposition 25.20, $P, Q, \mathrm{apex}$ — and hence $M := \mathrm{norm}(P + Q)$ and the corner's shared apex $v$ — do not depend on $\beta$ at this corner for $\beta \in \{A, B, C\}$, since proposition 25.20's exact $\pm 72^\circ$-rotation reduction for $N_{\mathrm{left}}, N_{\mathrm{right}}$ holds for every geodesic-built scheme regardless of any internal base choice. Only $R$ depends on $\beta$, and it depends on $\beta$ only through the binary distinction $\beta = A$ (*matched*: the formula's privileged vertex coincides with the corner's own near vertex) versus $\beta \in \{B, C\}$ (*mismatched*). Computed at $n = 64, \ldots, 2048$ and extrapolated (quadratic in $1/n$, four points), the resulting $\kappa_\infty$ is $1.013507$ for $\beta = A$ and $0.927702$ for $\beta = B$ and, separately, $0.927702$ for $\beta = C$ (agreement to the sixth decimal place between the two mismatched choices).*

*Verification.* A new diagnostic script (`sym_geo_privilege_isolation.py`) implements definition 25.24 for all three choices of $\beta$, evaluates the $A$-anchored corner's defect under each, and separately confirms the raw displacement match: comparing $d_{\mathrm{SYM}}(n)$ against $d_{\beta=B}(n)$ — both now defined at the identical label $(n{-}2, 1, 1)$, with no appeal to icosahedral symmetry to relate two different corners — gives ratio $1.00119530, 1.00030064, 1.00007527, 1.00001882, 1.00000446$ at $n = 256, 1024, 4096, 16384, 65536$, deviation from 1 shrinking by very close to a factor of 4 at each quadrupling: the same clean $O(1/n)$ convergence proposition 25.23 found across corners, now reproduced at a single, fixed corner. Against the matched case ($\beta = A$) at the same label, the ratio instead sits at a stable $0.78846, 0.78705, 0.78669, 0.78660, 0.78658$ — converging to proposition 25.23's previously unexplained constant $0.78658$, now recovered from a same-label comparison rather than a cross-corner one. □

**Key Result**

*[SYM's agreement with GEO's non-privileged role is not really about GEO's roles at all] Proposition 25.25 replaces section 25.7.13's open comparison between two constructions at two corners with a controlled experiment inside a single construction at one corner: hold the corner fixed, vary only whether the two-`slerp` formula's arbitrary choice of outer-`slerp` base happens to be the vertex that corner is geometrically near. That one bit — matched or mismatched — reproduces* both *numbers already on record (*1.013507 *and* 0.927702*) exactly, with no reference to which of the three vertices is involved or to any cross-corner symmetry. Since definition 25.21's spherical Fréchet mean has, by its own permutation-symmetric construction, no privileged base vertex at any corner whatsoever, SYM sits in the mismatched class* everywhere*, not by coincidence but structurally: there is no label at which SYM could land in the matched class, because the matched class requires exactly the kind of vertex-dependent formula asymmetry SYM is built to not have. GEO's own role A is accordingly not one of two comparable "roles" on equal footing with role B — it is the single exceptional case, among the three possible base choices at any corner, where an otherwise arbitrary formula convention happens to align with geometry; every other combination of corner and base choice, GEO's own included, falls in the other class together with SYM. This is the algebraic mechanism sections 25.7.12 and 25.7.13's todoboxes asked for, stated at the level of* which class *a construction falls into and why; it is not a closed-form derivation of either* 1.013507 *or* 0.927702 *themselves, nor of the $O(1/n)$ rate's coefficient, both of which remain open below.*

**Todo — verify primary source**

Two items are closed by proposition 25.25 and section 25.7.14 above: the coincidence is no longer cross-corner-symmetry-dependent (it is now exhibited at one fixed label), and it is no longer specific to GEO-versus-SYM (it is a property of the matched/mismatched dichotomy itself, which any future two-`slerp`-style or Karcher-mean-style scheme will fall into on the same terms). What remains genuinely open is narrower and quantitative rather than structural: a closed-form value for either $\kappa_\infty^{\text{matched}} = 1.013507$ or $\kappa_\infty^{\text{mismatched}} = 0.927702$ individually (neither has been derived from first principles anywhere in this chapter, only computed), and a first-order perturbative derivation of the $O(1/n)$ convergence rate's actual coefficient from definition 25.24's explicit formula, which proposition 25.25's verification exhibits numerically but does not derive. Section 25.7.8's separate, still-open RMS/GEO $\sim 3\%$ gap is unaffected by this subsection: RMS is built from recursive midpoint bisection, not any two-`slerp` formula, so the matched/mismatched dichotomy identified here does not directly apply to it, and nothing above narrows that gap further.

### 25.7.15 High-precision values, and a genuine parity asymmetry between the two classes

Proposition 25.25's two constants were computed at ordinary double precision and extrapolated only to six significant figures. This subsection recomputes both at arbitrary precision (60 decimal digits, exact algebraic icosahedron coordinates via $\varphi = (1+\sqrt{5})/2$) to test whether more precision reveals a recognizable closed form, and along the way finds a genuine structural difference between the two classes that the six-figure values could not show.

**Proposition 25.26** (The matched series contains only even powers of $1/n$; the mismatched series does not)**.** *Write $\kappa_\beta(n) = \kappa_\infty^\beta + c_1^\beta/n + c_2^\beta/n^2 + c_3^\beta/n^3 + \cdots$ for $\beta \in \{\text{matched}, \text{mismatched}\}$. Fitting an exact degree-7 polynomial in $x = 1/n$ through eight values of $\kappa_\beta(n)$ computed at 60-digit precision for $n = 200, 400, 800, \ldots, 25600$ (a genuine interpolation, not a least-squares approximation, since the number of points equals the number of fitted coefficients) gives, for the matched case, $c_1^{matched} = 1.4\times10^{-25}$ and $c_3^{matched} = 1.8\times10^{-17}$, both zero to the full precision computed, while $c_2^{matched} = -0.0064975165\ldots$ is genuinely nonzero. For the mismatched case, by contrast, every coefficient tested is nonzero: $c_1^{mismatched} = 0.0953901709\ldots$, $c_2^{mismatched} = 0.1289106495\ldots$, $c_3^{mismatched} = -0.0063566704\ldots$. The*

*matched construction's approach to its limit is a function of $1/n^2$ alone; the mismatched construction's is not.*

*Verification.* A new high-precision script (`highprec_kappa.py`, using `mpmath` at 60 decimal digits, with icosahedron vertices built exactly from $\varphi$ rather than floating-point square roots) computes $\kappa_\beta(n)$ for $n$ up to 25600 and solves the resulting Vandermonde system exactly for the polynomial coefficients in $x = 1/n$. Both extrapolated constants are stable to all 20 digits reported under every tested choice of how many of the eight points are used (five through eight), confirming the fit is not an artifact of the particular point set: $\kappa_\infty^{\text{matched}} = 1.0135065410165624666\ldots$ and $\kappa_\infty^{\text{mismatched}} = 0.9277017681986128652\ldots$. □

**Key Result**

*[A closed form for either constant is not among the natural icosahedral candidates] Following propositions 25.37 and 25.43's own methodology, the two twenty-digit values above are screened against some two dozen expressions natural to icosahedral geometry: powers and simple rational functions of $\varphi$ and $\sqrt{5}$, $\cos 36°, \cos 72°, \sin 54°$, $\arccos(1/\sqrt{5})$ (the icosahedron's own vertex-to-vertex central angle) in both radian and degree form, and simple ratios and products of the two constants with each other. None matches either constant to within $1 \times 10^{-6}$, let alone to the twenty digits computed. This is a genuine negative result in the same sense proposition 25.37's original screen was: it does not prove no closed form exists, but it rules out the first, most natural place to look, consistent with proposition 25.26's finding that the two constants are not simple functions of each other either (product 0.940 231 810, ratio 1.092 491 764, neither recognizable).*

**Todo — verify primary source**

Proposition 25.26's parity finding is new to this pass and is itself left unexplained: no argument is given for why the matched construction's local geometry should be even in $1/n$ (equivalently, symmetric under some formal $n \to -n$ or reflection operation on the construction) while the mismatched one is not, though it is consistent with, and may share a common source with, section 25.7.11's much earlier observation that role A's own displacement statistics behave more simply than role B's. Pending that explanation, the closed-form question for both constants stands exactly as section 25.7.14 left it, now with twenty correct digits to check any future candidate against instead of six.

### 25.7.16 The parity puzzle, reduced: one exact lemma, one exact symmetry, one honest unknown

Proposition 25.26's parity finding was left unexplained. This subsection does not close it, but reduces it: decomposing the matched-case defect into its three constituent angles shows that two of the three pieces are handled by exact, elementary facts, leaving a single scalar quantity carrying the entire open question.

**Lemma 25.27** (Rotation-image angle is independent of radius)**.** *Let $v$ be any point on the unit sphere, $a$ any point (the "axis"), and $\theta \in (0°, 180°)$. Write $N_\mp = \mathrm{Rot}_a(\mp\theta)\, v$ for $v$ rotated by $\mp\theta$ about the axis through $a$. Then*

$$\angle(N_-, v, N_+) = 180° - \theta,$$

*independent of the angular distance between $v$ and $a$.*

*Verification.* `pent_identity_check.py` confirms this at colatitudes $\rho = 1°, 10°, 30°, 60°, 89°, 120°, 150°$ (fixed $\theta = 72°$) and at $\theta = 10°, 30°, 45°, 60°, 90°, 120°$ (fixed $\rho = 37°$), all to 30-digit precision with zero residual. A synthetic proof: $v, N_-, N_+$ lie on a common circle of colatitude $\rho$ about $a$, symmetric under the reflection swapping $N_- \leftrightarrow N_+$ and fixing $v$; this reflection makes triangle $(N_-, v, N_+)$ isosceles with the two

base points separated by arc $2\theta$ along that circle, and the resulting apex angle at $v$ is a function of $\theta$ alone by a standard spherical-isosceles argument, independent of $\rho$. □

**Corollary 25.28** (`pent` is not asymptotically constant – it is exactly constant)**.** *Applying lemma 25.27 with $\theta = 72°$ to the role-A construction's $N_{\text{left}}, N_{\text{right}}$ (which are, by proposition 25.20, exactly $v$ rotated by $\mp 72°$ about the axis through $A$, for any $n$) gives* $\text{pent}(n) = \angle(N_{\text{left}}, v, N_{\text{right}}) \equiv 108°$ *exactly, for every $n$ – not a series in $1/n$ with vanishing higher coefficients, but a genuine constant. Section 25.7.15's finding that the matched series has $c_1 = c_3 = 0$ to 60-digit precision is, for this one term, not a numerical coincidence to be explained: it is a consequence of lemma 25.27 together with the $\mp 72°$ rotation identity already established.*

**Proposition 25.29** (`hex1` and `hex2` coincide exactly, at every $n$)**.** *In the role-A (matched) construction,* $\text{hex1}(n) = \text{hex2}(n)$ *exactly, for every $n$.*

*Verification.* The role-A construction is manifestly symmetric under exchanging $B \leftrightarrow C$: $P \leftrightarrow Q$ under this exchange (both are single-edge `slerp`s from $A$), and $R = \text{slerp}(A, \text{slerp}(B, C, \frac{1}{2}), 2/n)$ is fixed by it, since $\text{slerp}(B, C, \frac{1}{2}) = \text{slerp}(C, B, \frac{1}{2})$ exactly. Hence $N_{\text{int}} = \text{norm}(P + Q + R)$ is fixed by the exchange, while $N_{\text{left}} \leftrightarrow N_{\text{right}}$ swap under it (they are $v$ rotated by $\mp 72°$ about the same fixed axis). Since $\text{hex1} = \angle(N_{\text{right}}, v, N_{\text{int}})$ and $\text{hex2} = \angle(N_{\text{int}}, v, N_{\text{left}})$, the exchange maps one to the other. `angle_decomposition.py` confirms $\text{hex1}(n) - \text{hex2}(n)$ vanishes to the full 60-digit precision computed at $n = 200, 1600, 12800$ (residuals $\leq 7.2 \times 10^{-56}$, i.e. exactly zero within the arithmetic's own floor). □

**Key Result**

*[The parity question, reduced to a single scalar] Combining corollary 25.28 and proposition 25.29, the matched-case defect is*

$$\delta_{matched}(n) = 360° - 108° - 2\,\text{hex1}(n) = 252° - 2\,\text{hex1}(n).$$

*Proposition 25.26's finding that $\delta_{matched}$'s series in $1/n$ has $c_1 = c_3 = 0$ is therefore neither a three-way conspiracy among `pent`, `hex1`, `hex2`, nor a property of the angle-defect construction as a whole: `pent` contributes no series at all (corollary 25.28), and `hex1`, `hex2` are one and the same function (proposition 25.29), so no cancellation between distinct functions is occurring. The entire open question is whether* $\text{hex1}(n)$ *– a single angle, $\angle(N_{\text{right}}, v, N_{\text{int}})$, built from the two points ($v$, radial from $A$ at parameter $O(1/n)$; $N_{\text{int}}$, involving $R$) that are* not *related to each other by any exact symmetry already in hand – itself has a series in $1/n$ containing only even powers. `angle_decomposition.py`'s own Vandermonde fit on* $\text{hex1}(n)$ *alone gives odd coefficients $c_1, c_3, c_5 = O(10^{-23})$, i.e. zero within the fit's own numerical floor at this precision – consistent with, but not a proof of, that parity.*

### 25.7.17 The parity, proved: a 180° rotation, not an expansion

Section 25.7.16 left $\text{hex1}(n)$'s own parity as a numerical observation. This subsection proves it, and the proof is exact – an identity at every $t$, not an asymptotic statement about a Taylor series – and requires no curvature expansion at all.

**Definition 25.30** (The continuous extension of the role-A construction)**.** Fix an orthonormal tangent frame at $A$: let $\alpha = \angle(A, B) = \angle(A, C)$ (equal by the face's symmetry), $e_B = (B - \cos\alpha\, A)/\sin\alpha$, $e_C = (C - \cos\alpha\, A)/\sin\alpha$ the unit tangent directions at $A$ toward $B, C$; let $M = \text{slerp}(B, C, \frac{1}{2})$, $\beta = \angle(A, M)$, $e_M = (M - \cos\beta\, A)/\sin\beta$. For a *continuous*

real parameter $t$ (not restricted to $t = 1/n$ for integer $n$), define

$$P(t) = \cos(\alpha t)A + \sin(\alpha t)e_B,$$
$$Q(t) = \cos(\alpha t)A + \sin(\alpha t)e_C,$$
$$R(t) = \cos(2\beta t)A + \sin(2\beta t)e_M,$$

and $v(t) = \mathrm{norm}(A+P(t)+Q(t))$, $N_{\mathrm{int}}(t) = \mathrm{norm}(P(t)+Q(t)+R(t))$, $N_{\mathrm{right}}(t) = \mathrm{Rot}_A(72°)v(t)$, $N_{\mathrm{left}}(t) = \mathrm{Rot}_A(-72°)v(t)$. At $t = 1/n$ this reproduces definition 25.30's discrete role-A construction exactly, since $P(1/n) = \mathrm{slerp}(A, B, 1/n)$ etc. by the definition of `slerp` as $\cos(\theta t)$-$\sin(\theta t)$ motion along a geodesic at unit angular speed.

**Lemma 25.31** ($\mathrm{Rot}_A(180°)$ reverses only the tangential part)**.** *For any $\theta$ and any unit tangent vector $e$ at $A$ (i.e. $e \perp A$),* $\mathrm{Rot}_A(180°)\big[\cos\theta\, A + \sin\theta\, e\big] = \cos\theta\, A - \sin\theta\, e$.

*Proof.* Direct from Rodrigues' rotation formula: $\mathrm{Rot}_A(180°)X = X\cos 180° + (A \times X)\sin 180° + A(A \cdot X)(1 - \cos 180°) = -X + 2A(A \cdot X)$. With $X = \cos\theta\, A + \sin\theta\, e$ and $A \cdot X = \cos\theta$ (since $A \cdot e = 0$), this gives $-\cos\theta\, A - \sin\theta\, e + 2\cos\theta\, A = \cos\theta\, A - \sin\theta\, e$. □

**Proposition 25.32** (Exact reflection symmetry of the role-A construction)**.** *Write $\rho = \mathrm{Rot}_A(180°)$. Then, for every $t$:*

$$P(-t) = \rho[P(t)], \qquad Q(-t) = \rho[Q(t)], \qquad R(-t) = \rho[R(t)],$$

*and consequently $v(-t) = \rho[v(t)]$, $N_{\mathrm{int}}(-t) = \rho[N_{\mathrm{int}}(t)]$, $N_{\mathrm{right}}(-t) = \rho[N_{\mathrm{right}}(t)]$, $N_{\mathrm{left}}(-t) = \rho[N_{\mathrm{left}}(t)]$.*

*Proof.* The first three identities are lemma 25.31 applied to $P, Q, R$'s defining formulas (note cos is even and sin is odd in $t$, so $P(-t) = \cos(\alpha t)A - \sin(\alpha t)e_B$, exactly the right-hand side of lemma 25.31). Since $\rho$ is linear and fixes $A$, and $\mathrm{norm}(\cdot)$ commutes with any orthogonal linear map, $v(-t) = \mathrm{norm}(A+P(-t)+Q(-t)) = \mathrm{norm}(\rho[A+P(t)+Q(t)]) = \rho[v(t)]$, and likewise for $N_{\mathrm{int}}$. For $N_{\mathrm{right}}$: rotations about the same axis commute, so $N_{\mathrm{right}}(-t) = \mathrm{Rot}_A(72°)\rho[v(t)] = \rho\,\mathrm{Rot}_A(72°)[v(t)] = \rho[N_{\mathrm{right}}(t)]$, and symmetrically for $N_{\mathrm{left}}$. □

**Key Result**

*[*pent*, hex1, hex2, and the total defect are exactly even functions of t] Since $\rho$ is a rotation, it preserves all angles: for any three points $X, Y, Z$, $\angle(\rho X, \rho Y, \rho Z) = \angle(X, Y, Z)$. Proposition 25.32 says the entire configuration at $-t$ is the $\rho$-image of the configuration at $t$, so every angle built from it is unchanged:*

$$\mathrm{pent}(-t) = \mathrm{pent}(t), \qquad \mathrm{hex1}(-t) = \mathrm{hex1}(t), \qquad \mathrm{hex2}(-t) = \mathrm{hex2}(t),$$

*and hence $\delta_{matched}(-t) = \delta_{matched}(t)$. This is an* identity*, verified at $t = 0.2, 0.1, 0.05, 0.01, 0.001$ against $t = -0.2, -0.05$ etc. to* 20 *digits (`hex1_parity_test.py`), not merely a vanishing Taylor coefficient. It follows immediately that* hex1$(t)$*'s (and* pent*'s, and the total defect's) Taylor expansion about $t = 0$ contains only even powers of $t$, to* every *order – not just the seven coefficients proposition 25.26 checked. Section 25.7.16's open question is closed:* hex1$(n)$*'s parity was never a separate curvature fact requiring its own expansion; it is the same* 180°*-rotation symmetry that already gave corollary 25.28 its exact constant, applied to the whole configuration rather than to* pent*'s two points alone.*

*Remark* 25.33 (Why this is the general matched/mismatched dichotomy, not a feature of this one example). Nothing in proposition 25.32's proof used the specific values $\alpha, \beta$, the icosahedron, or the number 72°: it only used that every point entering the construction is $\cos(c\theta t)A + \sin(c\theta t)e$ for some fixed direction $e$ and fixed constant $c\theta$, and that any further points are built from these by rotations about the axis through $A$. This is exactly the

definition of the *matched* class from section 25.7.14: a base-privileged two-`slerp` construction is matched at a corner precisely when every point used to compute the defect there is $\mathrm{slerp}(A, \cdot, O(t))$-radial from that corner's own vertex $A$. Exact parity in $t$ is therefore not a property discovered empirically in this one construction; it is a structural consequence of matchedness itself, and will hold for any future two-`slerp`-style or Karcher-mean-style scheme's matched roles on the same terms. The mismatched class fails this by definition – its $R$ is $\mathrm{slerp}(B, \cdot, \cdot)$-based, not $A$-radial at small parameter – so proposition 25.32's hypothesis simply does not hold there, and proposition 25.26's nonzero odd coefficients for the mismatched series are not a separate fact requiring their own explanation; they are the generic case, absent any reason to expect otherwise.

**Todo — verify primary source**

What remains open after section 25.7.17 and remark 25.33 is narrower than before, on both fronts section 25.7.15 originally posed. First, the parity question itself is now fully closed for the matched class in general, not just for this one corner – but this proof says nothing about the *coefficients* of the surviving even powers $(c_2, c_4, \ldots)$, which are exactly what a closed-form value for $\kappa_\infty^{\text{matched}}$ would require, and remain purely numerical. Second, the mismatched class's own coefficients, all nonzero, still have no structural explanation beyond "the reflection argument does not apply" – a proof that a symmetry is absent is not itself a derivation of what replaces it, and section 25.7.11's $O(1/n)$-rate coefficient for either class is still read off finite differences, not derived from definitions 25.24 and 25.30's explicit formulas by any expansion carried out in this chapter. Section 25.7.8's RMS/GEO $\sim 3\%$ gap remains, as before, a separate item this mechanism does not touch.

### 25.7.18 Taking the reflection trick to RMS: two identities transplant exactly, the parity does not

Section 25.7.17's closing todobox left one candidate direction untried: RMS (recursive midpoint bisection) has its own, different kind of permutation symmetry – section 25.7.8 already described it informally as "manifestly symmetric under permuting a face's three vertices" – and it is natural to ask whether the same pent $\equiv 108°$, hex1 = hex2, full-parity-in-$1/n$ package that corollary 25.28, proposition 25.29, and section 25.7.17 built for GEO's matched class transplants onto RMS. It transplants exactly for two of the three pieces, and demonstrably does not for the third – and the reason it does not is itself informative about what made the matched-class argument work.

**Proposition 25.34** (RMS's label→point map is exactly invariant under all six permutations of $(A, B, C)$)**.** *Let* $\mathrm{rms}_{(A,B,C)}(i, j, k)$ *denote the point at label* $(i, j, k)$ *produced by section 25.7.9's recursive-bisection construction, seeded from an ordered triple* $(A, B, C)$*. For every permutation* $\sigma$ *of* $(A, B, C)$ *and every valid label* $(i, j, k)$*,*

$$\mathrm{rms}_{(A,B,C)}(i, j, k) = \mathrm{rms}_{\sigma(A,B,C)}\big(\sigma^{-1}(i, j, k)\big),$$

*i.e. the underlying point depends only on the unordered assignment of weights to vertices, never on which vertex the recursion happened to start from or which order its arguments were passed in.*

*Proof.* By structural induction on the recursion depth, it suffices to show invariance under (a) the 3-cycle $(A, B, C) \to (B, C, A)$ and (b) the transposition $(A, B, C) \to (A, C, B)$, since these generate all of $S_3$. Write $\mathrm{rec}(P_1, P_2, P_3, l_1, l_2, l_3, \ell)$ for one recursive call (as in section 25.7.9's construction): it computes $M_{12} = \mathrm{slerp}(P_1, P_2, \frac{1}{2})$, $M_{23} = \mathrm{slerp}(P_2, P_3, \frac{1}{2})$, $M_{31} = \mathrm{slerp}(P_3, P_1, \frac{1}{2})$ and the componentwise label midpoints $\mathrm{im}_{12}, \mathrm{im}_{23}, \mathrm{im}_{31}$, then recurses into four sub-calls built from $(P_1, M_{12}, M_{31})$, $(M_{12}, P_2, M_{23})$, $(M_{31}, M_{23}, P_3)$, $(M_{12}, M_{23}, M_{31})$

with matching labels. Both facts used below hold because $\mathrm{slerp}(p, q, \frac{1}{2}) = \mathrm{slerp}(q, p, \frac{1}{2}) = (p+q)/\|p+q\|$ is manifestly symmetric in its two arguments (already noted in section 25.7.9's own definition of RMS), and because componentwise averaging commutes with any fixed permutation of coordinate positions.

*Base case* ($\ell = 0$): trivial, since rec simply deposits $P_1, P_2, P_3$ at $l_1, l_2, l_3$, and both the 3-cycle and the transposition act as a relabelling of which point is deposited at which (permuted) label.

*Cyclic step*: relabelling $(P_1, P_2, P_3, l_1, l_2, l_3) \to (P_2, P_3, P_1, l_2, l_3, l_1)$ gives $M'_{12} = \mathrm{slerp}(P_2, P_3, \frac{1}{2}) = M_{23}$, $M'_{23} = \mathrm{slerp}(P_3, P_1, \frac{1}{2}) = M_{31}$, $M'_{31} = \mathrm{slerp}(P_1, P_2, \frac{1}{2}) = M_{12}$, and correspondingly $\mathrm{im}'_{12} = \mathrm{im}_{23}$, $\mathrm{im}'_{23} = \mathrm{im}_{31}$, $\mathrm{im}'_{31} = \mathrm{im}_{12}$; each of the four resulting sub-calls is then exactly a cyclic relabelling of one of the original four, so the inductive hypothesis at depth $\ell - 1$ closes the step.

*Transposition step*: relabelling $(P_1, P_2, P_3, l_1, l_2, l_3) \to (P_1, P_3, P_2, \sigma l_1, \sigma l_3, \sigma l_2)$ (where $\sigma$ swaps a label's 2nd and 3rd coordinate) gives $M'_{12} = \mathrm{slerp}(P_1, P_3, \frac{1}{2}) = M_{31}$, $M'_{23} = \mathrm{slerp}(P_3, P_2, \frac{1}{2}) = M_{23}$, $M'_{31} = \mathrm{slerp}(P_2, P_1, \frac{1}{2}) = M_{12}$, with labels $\mathrm{im}'_{12} = \sigma(\mathrm{im}_{31})$, $\mathrm{im}'_{23} = \sigma(\mathrm{im}_{23})$, $\mathrm{im}'_{31} = \sigma(\mathrm{im}_{12})$; matching each of the four resulting sub-calls against the original four (the middle sub-call matches only after also invoking the already-established cyclic invariance) again closes the step by the inductive hypothesis. □

**Epistemic Notice**

`scripts/rms_permutation_exact.py` checks proposition 25.34 directly, computationally, rather than resting on the induction alone: it builds the full RMS grid on one icosahedron face independently from all six orderings of $(A, B, C)$ at $n = 32$ and compares every one of the 561 resulting points against the corresponding relabelled point from the $(A, B, C)$-seeded grid. All six permutations agree with the base ordering to a maximum residual of exactly 0.000 – not merely small, but bit-identical, which is what the induction predicts: both sides of every identity above are literally the same sequence of `slerp` calls with arguments relabelled, so floating-point rounding cannot distinguish them.

**Key Result**

*[`pent` and `hex1`=`hex2` transplant onto RMS unchanged] Two of section 25.7.16's three pieces now apply to RMS directly. First,* $\mathrm{pent}_{\mathrm{RMS}}(n) \equiv 108°$ *exactly, for the same reason as corollary 25.28: proposition 25.20 already established* $N_{\mathrm{left}}, N_{\mathrm{right}} = \mathrm{Rot}_A(\mp 72°)v$ *for* any *scheme built from geodesics, RMS included, and lemma 25.27 turns that alone into the constant* 108° *with no further RMS-specific argument needed. Second,* $\mathrm{hex1}_{\mathrm{RMS}}(n) = \mathrm{hex2}_{\mathrm{RMS}}(n)$ *exactly, by the same argument as proposition 25.29: swapping* $B, C$ *fixes* $R$ *at label* $(n-2, 1, 1)$ *(symmetric in its last two coordinates) and hence* $N_{\mathrm{int}}$*, swaps* $P \leftrightarrow Q$ *and hence* $N_{\mathrm{left}} \leftrightarrow N_{\mathrm{right}}$*, and is an exact symmetry of RMS's own construction by proposition 25.34 rather than of a two-`slerp` formula's B-C symmetry. Both identities are confirmed computationally by `scripts/rms_pent_hex_parity.py` at* $n = 8, 16, \ldots, 1024$*:* $\mathrm{pent}_{\mathrm{RMS}}(n) = 108.0000000000°$ *at every level tested, and* $\mathrm{hex1}_{\mathrm{RMS}}(n) - \mathrm{hex2}_{\mathrm{RMS}}(n)$ *stays between* $1.1 \times 10^{-13}$ *and* $1.1 \times 10^{-11}$ *across the whole range with no growth trend – consistent with the accumulated float64 rounding floor of the* $O(\log_2 n)$ *sequential `slerp`/rotation calls involved, not with a genuine nonzero value. Section 25.7.16's reduction therefore holds for RMS unchanged in form:* $\delta_{\mathrm{RMS}}(n) = 252° - 2\,\mathrm{hex1}_{\mathrm{RMS}}(n)$.

**Key Result**

*[RMS's own local $\kappa_\infty$, by this same four-point reduction, lands within $5 \times 10^{-9}$ of the mismatched/SYM constant] Extrapolating $\kappa_{\mathrm{RMS}}(n)$ from proposition 25.20's local reduction the same way as every other construction in this subsection (Vandermonde fit in $x = 1/n$, degree 7, $n = 8$ to 1024, stable to the digits reported under every truncation from degree 2 up) gives $\kappa_\infty^{\mathrm{RMS,\,local}} = 0.92770177262$ – consistent with, and sharpening, section 25.7.11's earlier six-figure estimate ($\approx 0.9277$) of the same quantity. Set against proposition 25.26's 60-digit $\kappa_\infty^{\mathrm{mismatched}} = 0.9277017681986128652\ldots$, the two values differ by $4.4 \times 10^{-9}$ – roughly three orders of magnitude closer than section 25.7.12's own SYM/mismatched agreement ($4.77 \times 10^{-6}$) was, itself already flagged there as "coincidence-worth-explaining." Three constructions built by three genuinely different mechanisms – GEO's own non-privileged two-`slerp` role, SYM's global spherical Fréchet mean, and now RMS's recursive midpoint bisection – cluster on the same local pentagon-corner constant to a precision this chapter has no mechanism to explain for even one such pair, let alone three. This sharpens rather than resolves section 25.7.12's original open question; no new candidate explanation is proposed here.*

**Proposition 25.35** (Unlike the matched class, RMS's $\kappa(n)$ series shows a genuine, nonzero linear term)**.** *A degree-3 Vandermonde fit of $\kappa_{\mathrm{RMS}}(n)$ through the four points $n = 128, 256, 512, 1024$, evaluated* out of sample *at $n = 64$, predicts $\kappa_{\mathrm{RMS}}(64) = 0.92778932$ against the directly computed value $0.92779053$ – a residual of $1.2 \times 10^{-6}$. The even-powers-only model forced through the same four points (fitting $c_0 + c_2x^2 + c_4x^4 + c_6x^6$ only, with $c_1 = c_3 = 0$ imposed) predicts $0.92773612$ at the same held-out $n = 64$ – a residual of $5.4 \times 10^{-5}$, about 45 times larger. If RMS's series were genuinely even in $1/n$ the way GEO's matched class provably is, the even-only model should fit at least as well as the general one; it does not, by a wide and consistent margin.*

*Verification.* `scripts/rms_pent_hex_parity.py` performs this holdout comparison directly (both fits solved exactly via `mpmath` to eliminate solver round-off as a confound, leaving only the double-precision input data itself as the limiting factor); the same script also reports the standard (non-holdout) Vandermonde fit's own linear coefficient $c_1^{\mathrm{RMS}}$, which stabilizes to $-5.5 \times 10^{-6}$ across every range and degree tried from $n = [256, 1024]$ up to $n = [8, 1024]$ – small compared to $c_2^{\mathrm{RMS}} \approx 0.366$, but consistently signed and of a magnitude the holdout test independently confirms is real rather than fit noise. □

**Epistemic Notice**

This is not the same standard of proof as section 25.7.17's exact $t \to -t$ identity for GEO's matched class, and it cannot be, for a structural reason rather than a computational one: definition 25.30's argument needed a genuine continuous-$t$ closed formula $(P(t), Q(t), R(t)$ all literally $\cos(c\theta t)A + \sin(c\theta t)e)$ to apply $\mathrm{Rot}_A(180°)$ at an arbitrary real $t$ and read off parity to every order. RMS has no such formula – its grid exists only at $n = 2^{\mathrm{level}}$, built by a genuinely recursive, combinatorial process (proposition 25.34's own proof is by induction on recursion depth, not by evaluating a closed expression at a parameter), and no continuous extension of it to real $t$ is constructed or attempted here. Nor was an arbitrary-precision (e.g. `mpmath`) version of the RMS recursion run to confirm proposition 25.35's finding beyond double precision: at the levels needed for a Vandermonde fit sensitive to a $\sim 5 \times 10^{-6}$ linear coefficient, this would require building a full $n \times n$ grid recursively in arbitrary-precision arithmetic, which was not attempted on cost grounds. Proposition 25.35's holdout test is designed to be robust to this limitation (it compares two models' predictive residuals against each other, not either model's coefficients against zero in isolation), but it is evidence, not proof, of the kind section 25.7.17 supplied for the matched class.

**Todo — verify primary source**

The reflection trick, taken to RMS, closes two of three fronts and leaves the third open in a sharper, better-diagnosed form than before it was tried. Closed: RMS's permutation symmetry, previously only asserted informally, is now an exact, inductively proved theorem (proposition 25.34), and it is exactly this theorem – not RMS's construction in general – that reproduces pent $\equiv 108°$ and hex1 = hex2 (section 25.7.18) on the same terms as GEO's matched class. Open, and now understood structurally rather than merely unproved: full parity of $\kappa_{\mathrm{RMS}}(n)$'s series in $1/n$ does not hold, or at least the holdout evidence in proposition 25.35 points that way, and remark 25.33's own mechanism explains why the two closed identities were never enough to guarantee it – parity to *every* order needed proposition 25.32's continuous $t \to -t$ symmetry specifically, which is a strictly stronger statement than pointwise permutation-invariance at each discrete $n$, and RMS's recursive definition supplies the latter without obviously supplying (or excluding) the former. Whether some other, genuinely combinatorial substitute for the continuous reflection argument exists for recursively-defined grids – one that would settle proposition 25.35 the way section 25.7.17 settled the matched case, rather than only gathering holdout evidence against it – is not attempted here. Also newly open: section 25.7.18's three-way near-coincidence among RMS, SYM, and GEO's mismatched role has no proposed mechanism at all. Section 25.7.8's global RMS/GEO $\sim 3\%$ gap is untouched by any of this: the local reduction used throughout this subsection reproduces RMS's *local* four-point-reduction value, already known (section 25.7.11) to sit about 0.4% from RMS's independently, globally computed $\kappa_\infty = 0.923659$, and nothing here narrows that separate, pre-existing gap.

### 25.7.19 Closing the loop: does the role-mixture model actually predict the $\sim 3\%$ global gap?

Section 25.7.11 already contains everything needed to answer this, but never assembled the pieces into a single predicted number and checked it against the measured one; this subsection does exactly that. The model implicit in section 25.7.11 is: RMS carries no privileged corner at all (proposition 25.34's exact permutation symmetry), so it should behave, locally, like GEO's own non-privileged corner alone – $\kappa_\infty^{\mathrm{RMS,\,model}} := \kappa_\infty^{\text{role B=C}}$ – while GEO's global value should be the $(1,2)$-weighted mixture $\kappa_\infty^{\mathrm{GEO,\,model}} := (\kappa_\infty^{\text{role A}} + 2\kappa_\infty^{\text{role B=C}})/3$ already computed. Subtracting gives a closed-form *predicted gap*, $\kappa_\infty^{\mathrm{GEO,\,model}} - \kappa_\infty^{\mathrm{RMS,\,model}} = (\kappa_\infty^{\text{role A}} - \kappa_\infty^{\text{role B=C}})/3$, that never needs either scheme's own globally computed value as an input.

**Key Result**

*[The role-mixture model predicts the measured* $\sim 3\%$ *RMS/GEO gap to within about* $1\%$ *relative – and both schemes' models overshoot their own global value by almost exactly the same amount]* `scripts/kappa_gap_mechanism_quantified.py` *recomputes role A and role B to higher precision than section 25.7.11's original pass (extending the extrapolation range to* $n = 64$–$8192$ *and fitting cubic-, not only linear-, in-*$1/n$*), obtaining* $\kappa_\infty^{role\ A} = 1.01350655$, $\kappa_\infty^{role\ B=C} = 0.92770171$ *(both stable to the digits given between the quadratic and cubic fits, max residual* $\leq 8.8 \times 10^{-8}$*), and independently recomputes RMS's and GEO's own* globally *extrapolated* $\kappa_\infty$ *by full-mesh construction extended past section 25.7.8's original range (RMS to* $n = 128$*, GEO to* $n = 100$*, cubic-in-*$1/n$ *fit on the four largest frequencies in each case), obtaining* $\kappa_\infty^{\mathrm{RMS,\,global}} = 0.92365996$ *and* $\kappa_\infty^{\mathrm{GEO,\,global}} = 0.95197572$ *– consistent to five and four significant figures respectively with section 25.7.8's original* $0.923659$, $0.951982$*, and confirming this reimplementation is faithful. The four numbers combine as follows:*

$$\textit{model-predicted gap} = \frac{\kappa_\infty^{role\ A} - \kappa_\infty^{role\ B=C}}{3} = 0.02860161\ (2.9909\%\ \textit{relative}),$$

$$\textit{measured global gap} = \kappa_\infty^{\mathrm{GEO,\,global}} - \kappa_\infty^{\mathrm{RMS,\,global}} = 0.02831576\ (2.9744\%\ \textit{relative}).$$

*The model-to-actual ratio is* $1.0101$*: the purely local, role-based mechanism – built entirely from proposition 25.20's four-point reduction at a single representative face, with no reference to either scheme's full* $O(n^2)$ *mesh – predicts the measured global gap to within about* $1\%$ *relative of its own size, not merely the right sign and order of magnitude section 25.7.11 already established. This is the quantitative closure section 25.7.10's closing paragraph asked for: the* $\sim 3\%$ *RMS/GEO gap* is*, to* $99\%$ *accuracy of the gap's own magnitude, the effect of GEO mixing in one-third of a corner with R displaced* $27\%$ *further from the flat bisector than RMS's uniformly non-privileged corners ever are.*

*The residual is itself informative. Each model individually overshoots its own scheme's global value, and by nearly the same relative amount:* $\kappa_\infty^{\mathrm{RMS,\,model}}$ *exceeds* $\kappa_\infty^{\mathrm{RMS,\,global}}$ *by* $0.4376\%$*, and* $\kappa_\infty^{\mathrm{GEO,\,model}}$ *exceeds* $\kappa_\infty^{\mathrm{GEO,\,global}}$ *by* $0.4546\%$ *– a difference of only* $0.017$ *percentage points between the two overshoots, an order of magnitude smaller than either overshoot itself. Because the predicted* gap *is a difference of the two models, these two near-equal overshoots cancel almost exactly, which is the specific reason the gap prediction (*$1.01\times$*) comes out so much more accurate than either scheme's own model-vs-global agreement (*$1.0044\times$*,* $1.0045\times$*) already was on its own.*

**Epistemic Notice**

Section 25.7.19's near-coincidence between the RMS and GEO overshoots is not explained here. Proposition 25.20's reduction is exact for a single, isolated pentagon corner considered on its own two-face-plus-one-interior-triangle neighbourhood; it does not, by construction, account for any second-neighbour-shell effect by which a scheme's global mesh differs, even infinitesimally in the $n \to \infty$ limit, from the idealized single-corner picture the four-point reduction computes. That both RMS and GEO show almost exactly the same-sized overshoot is consistent with such a shared, scheme-independent second-order correction to the four-point reduction itself – the reduction is exact for $\delta(v)$ at any finite $n$ (proposition 25.20's own proof), so the overshoot is not an error in the reduction's algebra but a genuine feature of how the role-specific *extrapolated limit* of the reduction relates to the scheme's true global extrapolated limit – but this is a plausible diagnosis, not a derivation, and no mechanism producing specifically $\approx 0.44$–$0.45\%$ is proposed.

**Todo — verify primary source**

Section 25.7.19 answers section 25.7.10's closing question at essentially the level of accuracy the four-point reduction can offer: not an exact algebraic identity, but a 99%-accurate quantitative prediction of the $\sim 3\%$ RMS/GEO gap from purely local, role-based data. What remains open, sharpened rather than newly created by this subsection, is narrower than section 25.7.10 left it: (i) why the RMS-model and GEO-model overshoots of their own respective global values are so close to each other (0.4376% against 0.4546%) – whether this reflects one shared correction term or two unrelated $O(1/n)$-suppressed effects landing near the same size by chance; (ii) whether that common residual, once identified, could be folded into the four-point reduction to close the remaining 1% of the gap prediction exactly; and (iii), unchanged from section 25.7.20's own todobox, a closed-form value for $\kappa_\infty^{\text{role A}}$ or $\kappa_\infty^{\text{role B=C}}$ individually remains as open as $\kappa_\infty^{\text{SSB}}$'s own closed form – this subsection explains the *gap between* constructions quantitatively without explaining any one construction's value from first principles.

### 25.7.20 Toward a closed form for $\kappa_\infty$: a tighter numerical bound, and a negative screen

Proposition 25.18's todobox left two gaps: whether $\kappa(n)$ actually converges, and, if it does, why to 0.602 specifically rather than to some other constant between 0 and 1. This subsection does not close either gap, but sharpens the second one: it tightens the numerical estimate of $\kappa_\infty$ beyond proposition 25.18's linear-in-$1/n$ model, and then checks a list of the closed forms an icosahedrally symmetric construction might plausibly produce against that tighter estimate — finding, in every case tried, that they do not match.

**Proposition 25.36** (A quadratic-in-$1/n$ refit narrows $\kappa_\infty$ to five significant figures, and shifts it slightly)**.** *Refitting $\kappa(n)$ against $\kappa_\infty + c/n + d/n^2$ (one further term beyond proposition 25.18's model) by ordinary least squares: over the seven points $n = 10$ to 50, $\kappa_\infty = 0.602700$, $c = 0.9122$, $d = 0.4364$, with every residual under $2.1 \times 10^{-5}$; over the five points $n = 16$ to 50, $\kappa_\infty = 0.602782$, $c = 0.9076$, $d = 0.4910$, with every residual under $2.6 \times 10^{-6}$ — roughly an order of magnitude tighter than the corresponding linear fits' residuals at the same ranges (proposition 25.18).*

**Key Result**

*[The refined estimate is $\kappa_\infty \approx 0.6027$–$0.6028$, not 0.602 flat — and the extrapolated value is itself model-dependent at the fourth significant figure] Proposition 25.36's two quadratic fits agree with each other to five significant figures (0.602700 vs. 0.602782, a spread of $8\times10^{-5}$), tighter agreement than proposition 25.18's linear fits managed across their own restricted ranges (0.6015 to 0.6024, a spread of $9 \times 10^{-4}$) — consistent with the quadratic model absorbing curvature the linear model left as residual, and consistent with genuine convergence rather than continued drift. But the two models do not agree with* each other*: the linear fits' central estimate ($\approx 0.602$) and the quadratic fits' ($\approx 0.6027$–$0.6028$) differ by about $6 \times 10^{-4}$, several times either model's own internal spread. This is exactly the kind of model-dependence a purely numerical extrapolation cannot fully resolve on its own: both models fit the available data well by their own residuals, and choosing between them — or between either and a further, cubic-in-$1/n$ correction not attempted here — is not something table 25.6's thirteen points can settle. The honest summary is a numerical bound, not a numerical value: $\kappa_\infty \approx 0.6027 \pm 0.0005$, tighter than proposition 25.18's 0.602-to-three-figures but not yet pinned down to a fourth figure with confidence.*

**Proposition 25.37** (A screen of fifteen natural closed-form candidates, none within tolerance)**.** *Testing fifteen closed-form candidates built from the golden ratio $\varphi$, $\sqrt{5}$, $\pi$, and small rationals — quantities that arise naturally in icosahedral geometry, where $\varphi$ in particular is ubiquitous (icosahedron vertex coordinates, dihedral angles, the pentagon's own diagonal-to-side ratio) — against section 25.7.20's central estimate 0.6027 at a tolerance of $2 \times 10^{-4}$*

*(comparable to the two quadratic fits' own mutual spread), every candidate misses: the closest,* $1/\varphi = \varphi - 1 \approx 0.6180$*, misses by* $1.5 \times 10^{-2}$*, roughly seventy-five times the tolerance;* $3/5 = 0.6000$ *misses by* $2.7 \times 10^{-3}$*, about fourteen times the tolerance. A further search over rationals* $p/q$ *with* $q \leq 30$ *finds none within tolerance either.*

**Key Result**

*[Ruling out the obvious candidates is itself informative] Proposition 25.37's negative result narrows, rather than merely fails to answer, the closed-form question:* $\kappa_\infty$ *is not* $1/\varphi$ *(a natural first guess, given how often* $\varphi$ *appears in icosahedral coordinates, and close enough in magnitude —* 0.618 *vs.* 0.603 *— that it might have been mistaken for a match at proposition 25.18's original, looser precision), nor a simple rational, nor any of the handful of other icosahedrally-flavored combinations tried here. This is consistent with* $\kappa_\infty$ *being a genuine construction-specific constant rather than a universal icosahedral one: unlike* $108°$ *(section 25.7.1, a consequence of the regular pentagon's interior angle alone, true for any convention) or the exact* $36°/n^2$ *mean (section 25.7.7, a consequence of Gauss–Bonnet alone),* $\kappa_\infty$ *measures a property of* this specific *construction's discretization error — centroid-taking, radial renormalization, and SVD-plane angle measurement, composed — evaluated at one particular location (the pentagon corners) relative to the sphere's overall average, and there is no a priori reason a quantity defined that way should reduce to a simple algebraic number at all.*

**Todo — verify primary source**

What a genuine derivation of $\kappa_\infty$ would require, based on section 25.7.20's reasoning, is not a better numerical fit or a wider candidate search, but an asymptotic analysis of the limiting relative curvature-density profile $h(v)$ that section 25.7.6 already conjectures $\delta_n(v) \approx \kappa(n) \cdot h(v)$ converges to: specifically, the value of $h$ at (or approaching) an icosahedral-vertex direction, divided by $h$'s average over the sphere, which is exactly what $\kappa_\infty$ computes in the discrete setting. This chapter has not attempted that analysis — it would need a continuum limit of the barycentric-subdivision-plus-renormalization construction's discretization error near a five-fold symmetric point, as distinct from a generic point of the triangulated sphere, which is a genuinely different (and harder) problem than anything solved by propositions 25.7 and 25.11's exact global identities. Short of that analysis, section 25.7.20's bound $\kappa_\infty \approx 0.6027 \pm 0.0005$ is the most this chapter can offer, and proposition 25.37's negative screen — while only fifteen candidates, not exhaustive — is evidence against the value having a simple closed form of the kinds tried, not proof that no closed form exists. Section 26.10 should record this as a distinct open item from the convergence question itself: even granting that $\kappa(n) \to \kappa_\infty$ (still unproven), identifying $\kappa_\infty$ in closed form is a separate, and on present evidence harder, problem.

### 25.7.21 The right existing machinery for this question, located rather than built

A later pass, asked directly to attempt the continuum-limit analysis this section's own todobox describes, did not carry out that analysis from first principles — what follows explains why not, and identifies a specific pre-existing mathematical framework that is the natural tool for it, rather than leaving the todobox as a bare statement of difficulty.

The angle defect at a vertex of a triangulated surface, which is exactly the quantity $\delta_n(v)$ this chapter computes, is the curvature invariant of *Regge calculus*: a piecewise-flat approximation scheme for curved manifolds, in which curvature is concentrated (as a delta function, in the continuum reading) at each vertex, with magnitude equal to the angle defect there. The question this section asks — as a triangulated sphere is refined, does the vertex-by-vertex defect converge to a specific limiting density, and at what rate, and is the rate uniform across combinatorially different vertex types (fivefold pentagon corners versus generic sixfold vertices)? — is precisely the convergence question Regge calculus's own foundational literature addresses in general, not only for this construction.

**Key Result**

*[Two specific, precisely citable papers supply exactly the missing machinery, neither consulted in this pass] Two papers, both from 1984 and both still standard references in the Regge-calculus literature, prove general convergence results of the kind this section's todobox calls for. Cheeger, Müller, and Schrader, "On the curvature of piecewise flat spaces" (*Communications in Mathematical Physics *92 (1984), 405–454), prove that the Regge curvature* measure *of a sequence of triangulations refining a smooth Riemannian manifold converges weakly (as a measure, i.e. against smooth test functions) to the manifold's own smooth curvature measure — a rigorous, general statement of exactly the fact this chapter's proposition 25.7 already establishes in this construction's own special case (the total defect is exactly* 720° *at every finite $n$, matching $\int_{S^2} K\, dA = 4\pi$ up to units, with no limit required). Separately, and more directly relevant to $\kappa_\infty$ itself, Feinberg, Friedberg, Lee, and Ren, "Lattice gravity near the continuum limit" (*Nuclear Physics B *245 (1984), 343–368), prove a sharper, quantitative result for a general ("regular or irregular") triangulated lattice: the deviation of the discrete (Regge) curvature from the continuum value can be expanded as a power series in the square of the link length, with explicit general formulas given for the leading coefficient of that series. This is not a weak, measure-level statement but a pointwise perturbative expansion — exactly the order-by-order structure $\kappa(n) = \kappa_\infty + c/n + d/n^2 + \cdots$ that propositions 25.18 and 25.36 found numerically without a theoretical derivation of the coefficients. Neither paper has been obtained or read in this pass; both citations are drawn from their consistent appearance, with matching volume and page numbers, across five independent Regge-calculus review articles and reference lists consulted 1 Aug 2026 (a loop-quantum-gravity review, three general-relativity papers on Regge-calculus convergence and positivity, and a 2024 preprint on lattice curvature near the continuum limit that names the Feinberg–Friedberg–Lee–Ren result specifically as establishing $O(\ell^2)$ convergence of the discrete scalar curvature on smoothed lattices).*

**Key Result**

### 25.7.24 A candidate mechanism, tried directly: does the defect simply track a vertex-associated area?

Rather than continuing to pursue the Feinberg–Friedberg–Lee–Ren route, a later pass asked a more basic question first: independently of that paper's specific perturbative coefficients, is there an elementary geometric reason $\kappa(n)$ should approach a finite nonzero limit at all, testable without obtaining any further literature?

**Definition 25.38** (A vertex-associated area, in the same plane-projected convention already used for angles)**.** For every subdivision vertex $A$ of $\mathcal{P}_n$'s underlying triangulation (an original icosahedron vertex, giving a pentagonal Goldberg face, or a later subdivision vertex, giving a hexagonal one), let $\mathrm{Area}(A)$ be the planar area of the SVD-projected polygon already constructed for definition 25.10's angle convention — the same projection, reused for area rather than angle. For a Goldberg vertex $P$ (a small-triangle centroid, incident to exactly three such faces $A, B, C$), define

$$A(P) \;:=\; \frac{\mathrm{Area}(A)}{\mathrm{val}(A)} + \frac{\mathrm{Area}(B)}{\mathrm{val}(B)} + \frac{\mathrm{Area}(C)}{\mathrm{val}(C)},$$

where $\mathrm{val}(A) \in \{5, 6\}$ is the number of vertices of face $A$. This splits each face's area equally among its own corners, so that $\sum_P A(P) = \sum_A \mathrm{Area}(A)$ exactly, by construction, for every $n$.

The motivating hypothesis is a standard fact from discrete differential geometry, already implicit in section 25.7.21's description of the Cheeger–Müller–Schrader *measure*-level convergence theorem: if angle defect converges to Gaussian curvature as a measure, then at leading order, for a small enough patch, $\delta(v) \approx K \cdot A(v)$ for *every* vertex regardless of its combinatorial valence — the defect is governed by how much area is associated to that vertex, not by whether it happens to sit at a fivefold or sixfold combinatorial point. If this leading-order proportionality holds uniformly across both vertex types as $n \to \infty$, it predicts $\kappa(n) \to A_{\mathrm{pent}}(\infty)/\bar{A}(\infty)$: the same limit definition 25.17 defines for the defect ratio, but computed from areas, which are visible and computable without any curvature perturbation theory at all.

**Key Result**

*[The area ratio tracks $\kappa(n)$ closely at every tested frequency, and is itself far better behaved] A new, independent script (`scripts/kappa_area_hypothesis.py`, sharing only the icosahedron-construction and SVD-projection routines with `goldberg_defect.py`, not its defect computation) computes $\kappa_{\mathrm{area}}(n) := A_{\mathrm{pent}}(n)/\bar{A}(n)$ at the same thirteen frequencies $n = 2$ to 50 already tabulated for $\kappa(n)$ in table 25.6. The two ratios agree closely at every frequency tested: $\kappa_{\mathrm{area}}(2) = 0.984$ against $\kappa(2) = 1.011$; $\kappa_{\mathrm{area}}(10) = 0.687$ against $\kappa(10) = 0.698$; $\kappa_{\mathrm{area}}(50) = 0.600$ against $\kappa(50) = 0.621$ — a relative gap of 2.7 %, 1.6 %, and 3.4 % respectively, never exceeding 3.5 % across the full range, and the two sequences move together (both monotonically decreasing, both flattening at large $n$) rather than merely passing near each other at isolated points. As a byproduct, $\sum_P A(P)$ was checked against the unit sphere's exact area $4\pi$ at every $n$: the ratio rises from 0.9555 at $n = 2$ to 0.99993 at $n = 50$, the same qualitative approach to exactness section 25.7.2 already found for the total* defect *under this convention, now confirmed independently for total* area*.*

**Key Result**

*[The area ratio's own extrapolated limit is close to, but measurably below, $\kappa_\infty$] Extending the same computation to $n = 60, 70, 80, 100$ (feasible in under a minute per frequency on the hardware this chapter was developed on) and fitting $\kappa_{\text{area}}(n)$ against $\kappa_{\text{area},\infty} + c/n + d/n^2$ by ordinary least squares gives $\kappa_{\text{area},\infty} = 0.57976$, stable to five significant figures whether the fit is restricted to $n \geq 16$ or to $n \geq 30$ (residuals under $9 \times 10^{-6}$ and $5 \times 10^{-7}$ respectively) — an extrapolation at least as clean as proposition 25.36's own fit to the defect ratio directly, and requiring no assumption about which of propositions 25.18 and 25.36's two competing models for $\kappa(n)$ is correct. This value is close to but distinct from $\kappa_\infty \approx 0.6028$ (sections 25.7.20 and 25.7.28): the two differ by about 0.023, or roughly 4 % of $\kappa_\infty$'s own value — a gap an order of magnitude larger than either fit's own internal residuals, and therefore a real discrepancy between the two quantities, not fitting noise.*

**Todo — verify primary source**

Section 25.7.24 together give a mixed but informative answer to this subsection's opening question. The area-only mechanism is not merely plausible in principle; it is a derived (not fitted) geometric quantity that reproduces roughly 96 % of $\kappa_\infty$'s departure from 1, extrapolating cleanly to its own limit by the identical fitting procedure already applied to the defect ratio itself. It is not, however, the whole mechanism: a genuine, reproducible 4 % gap separates $\kappa_{\text{area},\infty}$ from $\kappa_\infty$, meaning $\delta(v) \propto A(v)$ alone — with a single, valence-independent constant of proportionality $K$ — is not exactly true even in the $n \to \infty$ limit of this construction; some further, valence-dependent correction survives. This is consistent with, rather than a refutation of, the picture section 25.7.21's todobox already anticipated: a leading-order Gauss–Bonnet argument (area alone) captures the dominant effect for free, without needing the Feinberg–Friedberg–Lee–Ren machinery at all, while the residual 4 % is exactly the kind of shape- or anisotropy-dependent correction that a genuine $O(\ell^2)$ expansion — sensitive to more than a vertex's dual-cell area, since it depends on the detailed local metric and mesh geometry, not area alone — would be needed to account for. What remains open is now sharper on both ends: not "why is $\kappa_\infty \approx 0.603$" as an undifferentiated question, but "why does a pure area argument already explain $\kappa_{\text{area},\infty} \approx 0.580$ of it, and what specific shape-dependent mechanism accounts for the remaining 0.023" — the second half being, plausibly, exactly where a specialization of the Feinberg–Friedberg–Lee–Ren formula would enter, were the paper obtained.

### 25.7.25 A second candidate, tried directly: does local shape explain the residual?

Section 25.7.24's todobox named a specific candidate for the residual 4 %: since angle defect is fundamentally an angular quantity, two vertex-associated cells of equal area but different elongation need not subtend equal defect, so a shape- or anisotropy-dependent correction to the pure area model is the natural next thing to try — and, unlike the Feinberg–Friedberg–Lee–Ren route, testable directly with the tools already in hand, without obtaining any further literature.

**Definition 25.39** (Cell eccentricity, in the same plane-projected convention already used for area and angle)**.** For every subdivision vertex $A$ with plane-projected polygon coordinates already constructed for definition 25.38 (the SVD-projected star of face centroids around $A$), let $\text{ecc}(A) := \sqrt{1 - \lambda_{\min}/\lambda_{\max}}$, where $\lambda_{\min} \leq \lambda_{\max}$ are the eigenvalues of the $2 \times 2$ covariance matrix of that polygon's own projected vertices. This is 0 for a perfectly isotropic (e.g. regular) polygon and approaches 1 for an increasingly elongated, needle-like one — a standard second-moment shape statistic, not specific to this construction. For a Goldberg vertex $P$ incident to faces $A, B, C$, define $\text{shape}(P)$ as the area-weighted average of $\text{ecc}(A), \text{ecc}(B), \text{ecc}(C)$, using definition 25.38's own per-corner area shares $\text{Area}(A)/\text{val}(A)$ as weights, so that the same partition of $A(P)$ underlies both quantities.

**Key Result**

*[A real correlation exists between local shape and local defect-per-area, but it weakens as $n$ grows]* *A new script (`scripts/kappa_shape_hypothesis.py`, reusing `kappa_area_hypothesis.py`'s construction and `goldberg_defect.py`'s plane-projected angle routine) computes, for every Goldberg vertex $P$ at a given frequency, the local ratio $\rho(P) := \delta(P)/A(P)$ — constant across all $P$ under the pure area hypothesis — alongside* $\mathrm{shape}(P)$*, and correlates the two across the full vertex population. At $n = 10$, $r = -0.56$ ($R^2 = 0.32$ for a linear fit of $\rho$ on shape); at $n = 40$, $r = -0.15$ ($R^2 = 0.02$). The correlation is real (consistently negative, well outside sampling noise at any of the eight frequencies $n = 4$ to 40 tested) but its strength drops by roughly a factor of four over this range, moving the wrong way for a mechanism meant to explain a residual that section 25.7.24 shows is* not *shrinking with $n$ over the same range.*

**Key Result**

*[Pentagon-corner cells carry zero shape variance to exploit; the correlation lives entirely within the hexagon population]* *Restricting to the sixty pentagon-corner vertices alone at any tested frequency, both $\rho(P)$ and* $\mathrm{shape}(P)$ *are constant to machine precision (standard deviation below $2 \times 10^{-9}$ at $n = 40$) — the same exact twelve-fold icosahedral symmetry section 25.7.1 already established for the raw pentagon-corner angle, now confirmed for cell shape as well. Restricting instead to the hexagon-only population, the correlation of section 25.7.25 persists and is, if anything, slightly stronger in isolation ($r = -0.59$ at $n = 10$, $r = -0.18$ at $n = 40$, tracking the mixed-population values closely) — but since the pentagon side of the ratio cannot move at all, this within-hexagon relationship has no lever with which to shift the pentagon/hexagon-mean gap that defines $\kappa(n)$ itself.*

**Key Result**

*[Fitting the correction in does not close the gap]* *Fitting $\rho(P) \approx K + K_2 \cdot \mathrm{shape}(P)$ by ordinary least squares and forming the model's own predicted $\kappa_{\mathrm{shape}}(n) := \overline{A \cdot (K + K_2 \cdot \mathrm{shape})}\big|_{\mathrm{pent}} / \overline{A \cdot (K + K_2 \cdot \mathrm{shape})}$ gives, at $n = 40$, $\kappa_{\mathrm{shape}}(40) = 0.605370$ against $\kappa_{\mathrm{area}}(40) = 0.605331$ and the true $\kappa(40) = 0.625777$: a shift of $4 \times 10^{-5}$ in the right direction, against a gap of $2.04 \times 10^{-2}$ — closing under 0.2 % of it. The same pattern holds at every frequency tested ($n = 4$ to 40): the fitted correction moves $\kappa_{\mathrm{shape}}(n)$ negligibly relative to $\kappa_{\mathrm{area}}(n)$, exactly as section 25.7.25 predicts it must, since the pentagon-side average the fit would need to move is already an exact constant with no shape variance for any linear correction to act on.*

**Todo — verify primary source**

Section 25.7.25 together give a clear negative verdict on this specific hypothesis, and a sharper one than a bare "did not work" would be: cell eccentricity, in the sense definition 25.39 makes precise, is not the residual's source, not because no relationship between shape and defect-per-area exists (section 25.7.25 show a real one, among hexagon cells specifically) but because that relationship cannot reach the quantity that needs explaining: section 25.7.25's machine-precision symmetry means the sixty pentagon-corner cells present a single, fixed shape value at every $n$, with no internal variance for a shape-based correction to leverage against the hexagon population's own average. Whatever separates the pentagon-corner ratio from the hexagon-mean ratio, in other words, is tied to being a fivefold-valent cell specifically — a discrete combinatorial fact — and not to any continuously-varying shape gradient of the kind ordinary hexagon-to-hexagon variation exhibits, which is exactly the quantity definition 25.39 measures and section 25.7.25 shows is real but fails to track the residual's own $n$-dependence. This narrows section 25.7.24's open question once more: not "what shape-dependent correction explains the residual," which section 25.7.25 answers in the negative for eccentricity specifically, but "what property specific to fivefold combinatorial valence, as opposed to a cell's continuous geometric shape, produces the residual 0.023" — a question closer in kind to section 25.7.21's literature on disclinations (which is, after all, about the topological fact of fivefold coordination itself, not about how round or elongated any particular disclinated cell happens to be) than to an ordinary shape-correction term. The Feinberg–Friedberg–Lee–Ren expansion, still unobtained (sections 25.7.22 and 25.7.23), remains the one identified route that could in principle settle this by construction rather than by further screening of ad hoc candidates.

### 25.7.26 A third candidate, tried directly: is the residual an artifact of measuring area on the flat, rather than the curved, faces?

Section 25.7.24's area ratio and section 25.7.25's shape correction both used definition 25.38's convention of measuring each Goldberg face's area after projecting it onto its own best-fit *plane* — a reasonable convention, but not obviously the one the leading-order Gauss–Bonnet heuristic ($\delta(v) \approx K \cdot A(v)$) actually calls for, since that heuristic is about area *on the sphere*, and section 25.7.1 already established that Goldberg hexagons are measurably non-planar. A natural, independently testable hypothesis is that some of the 0.023 residual (section 25.7.24) is simply an artifact of flattening a non-planar face before measuring it, rather than a genuine failure of the area heuristic itself — testable directly, without the Feinberg–Friedberg–Lee–Ren paper, by replacing the flat measurement with an exact one.

**Definition 25.40** (Exact spherical face area, via Girard's theorem)**.** For a Goldberg face with vertex set $\{c_1, \dots, c_m\}$ (each $c_i$ a triangle centroid already renormalized to the unit sphere, per definition 25.5), fan-triangulate from $c_1$ into the $m-2$ spherical triangles $(c_1, c_i, c_{i+1})$, $i = 2, \dots, m-1$, and sum each one's solid angle via the van Oosterom–Strackee formula

$$E(a,b,c) \;=\; \Big|\, 2\arctan_2\big(a\cdot(b\times c),\; 1 + a\cdot b + b\cdot c + c\cdot a\big)\Big|,$$

a standard, numerically stable closed form for the area of the spherical triangle with unit-vector vertices $a, b, c$, equal by Girard's theorem to its own angle-sum excess over $\pi$. Define $\mathrm{Area}_{\mathrm{sph}}(A) := \sum_{i=2}^{m-1} E(c_1, c_i, c_{i+1})$, and split it among the face's own corners by the identical $1/\mathrm{val}(A)$ rule definition 25.38 already used for the flat convention, giving a spherical-area analogue $A_{\mathrm{sph}}(P)$ and ratio $\kappa_{\mathrm{area,sph}}(n) := \overline{A_{\mathrm{sph}}}_{\,\mathrm{pent}}(n)/\overline{A_{\mathrm{sph}}}(n)$.

**Proposition 25.41** (Total spherical area is exact at every $n$, unlike the flat convention's)**.** *`scripts/kappa_girard_area.py` computes $\sum_A \mathrm{Area}_{\mathrm{sph}}(A)$ at $n = 2, 3, 4, 5, 6, 8, 10, 12, 16, 20, 30, 40, 50, 60, 70, 80, 100$ and finds it equal to $4\pi$ to within $5 \times 10^{-15}$ (floating-point roundoff) at every single frequency tested — not approaching $4\pi$ as $n$ grows, as section 25.7.24's flat total-area ratio does (*0.9555 *at $n = 2$ rising only to*

0.99993 *at* $n = 50$*), but equal to it exactly, at every* $n$ *including the coarsest. This is expected rather than surprising: the Goldberg faces genuinely partition the unit sphere by construction (definition 25.5), and definition 25.40's fan-triangulated Girard sum is, face by face, exactly the Gauss–Bonnet accounting for that partition; the check confirms the script's implementation rather than testing a live hypothesis.*

**Key Result**

*[The two area conventions agree closely, and extrapolate to the same limit] At every tested frequency,* $\kappa_{\text{area,sph}}(n)$ *and section 25.7.24's flat* $\kappa_{\text{area}}(n)$ *agree to four or five significant figures:* 0.983575 *against* 0.984172 *at* $n = 2$*;* 0.686457 *against* 0.686850 *at* $n = 10$*;* 0.600125 *against* 0.600143 *at* $n = 50$*. Fitting* $\kappa_{\text{area,sph}}(n)$ *against* $\kappa_\infty + c/n + d/n^2$ *over* $n \geq 16$ *(the identical procedure section 25.7.24 used) gives* $\kappa_{\text{area,sph},\infty} = 0.579735$*, and restricting to* $n \geq 30$ *gives* 0.579758 *— indistinguishable, to five significant figures, from section 25.7.24's flat-convention value* 0.57976*, and leaving the same roughly* 0.023 *gap against the true* $\kappa_\infty \approx 0.60282$ *(section 25.7.28) that the flat convention already left.*

**Todo — verify primary source**

This is a clean negative result, not an inconclusive one: replacing definition 25.38's flat, best-fit-plane area with definition 25.40's exact spherical area — a strictly more faithful measurement of "how much of the sphere's own surface belongs to this vertex," verified independently by proposition 25.41's exact total — changes the extrapolated area ratio by less than $3 \times 10^{-5}$, essentially nothing against the 0.023 gap it was tested against. The residual is therefore not, even in part, an artifact of definition 25.38's particular flattening convention: both a flat and a genuinely curved area measure agree with each other far more closely than either agrees with the true defect ratio $\kappa_\infty$. Combined with section 25.7.25's equally clean negative result for cell eccentricity, two of the most natural non-perturbative explanations for the gap — a measurement-convention artifact, and a continuous shape correction — are now both ruled out directly rather than left untested, sharpening section 25.7.25's own conclusion rather than merely repeating it: whatever accounts for the residual 0.023 is a genuine property of fivefold combinatorial valence under a chord-based (straight-edge) discretization specifically, not recoverable from any area or shape statistic computed on the vertices' own positions, curved or flat. This leaves the Feinberg–Friedberg–Lee–Ren route (sections 25.7.22 and 25.7.23), still unobtained, as the one remaining identified path that addresses the discretization's local perturbative structure directly rather than screening a further geometric statistic against it — a conclusion this pass reaches by elimination, having tried the two most obvious alternatives and found both cleanly insufficient, rather than by assumption.

### 25.7.27 Extending to $n = 100$: the model gap narrows, and $\kappa_\infty$ tightens to five figures

Section 25.7.20 left an unresolved model-dependence: the linear-in-$1/n$ and quadratic-in-$1/n$ fits, run over the same $n \leq 50$ data, extrapolated to visibly different values (0.602 vs. 0.6027–0.6028). This subsection extends the pentagon-corner defect sequence four frequencies further, to $n = 60, 70, 80, 100$ (computed with the identical plane-projected script as every other entry in this chapter; mesh construction and star-ordering cost grow with $n$, and $n = 100$ — 200,000 Goldberg vertices — takes under a minute on the hardware this chapter was developed on), and asks whether more data narrows the gap between the two models rather than just adding more points to each.

**Proposition 25.42** (Linear, quadratic, and cubic refits over ranges reaching $n = 100$)**.** *Refitting* $\kappa(n)$ *against* $1/n$ *(linear),* $\kappa_\infty + c/n + d/n^2$ *(quadratic), and* $\kappa_\infty + c/n + d/n^2 + e/n^3$ *(cubic), over ranges ending at the new* $n = 100$ *point: linearly,* $n = 30$ *to* 100 *gives* $\kappa_\infty = 0.602601$ *(max residual* $3.9 \times 10^{-5}$*),* $n = 50$*-to-*100 *gives* 0.602704 *(*$7.5 \times 10^{-6}$*),* $n = 60$*-to-*100

| $n$ | 60 | 70 | 80 | 100 |
|---|---|---|---|---|
| Pentagon-corner defect (°) | 0.0061805 | 0.0045247 | 0.0034550 | 0.0022029 |
| $\kappa(n)$ | 0.618049 | 0.615855 | 0.614214 | 0.611922 |

Table 25.7: Pentagon-corner defect and $\kappa(n)$, plane-projected convention, extending tables 25.4 and 25.6 to four further frequencies. The total defect sum remains 720.000000° to within $3 \times 10^{-9}$ degrees at all four, the same exactness proposition 25.11 established at $n = 2, 3, 4$ and section 25.7.7 guarantees at every $n$.

*gives* 0.602727 *(*$2.9 \times 10^{-6}$*); quadratically,* $n = 30$*-to-*100 *gives* 0.602819 *(*$3.4 \times 10^{-7}$*),* $n = 50$*-to-*100 *gives* 0.602822 *(*$2.3 \times 10^{-8}$*); cubically,* $n = 30$*-to-*100 *gives* 0.602824 *(*$4.1 \times 10^{-9}$*),* $n = 16$*-to-*100 *gives* 0.602825 *(*$1.7 \times 10^{-7}$*).*

**Key Result**

*[The gap between models closes as data are added, exactly as genuine convergence — rather than an artifact of either model — would predict] Three things change together as* $n = 60, 70, 80, 100$ *are added, and all three point the same way. First, each model's own internal spread shrinks: the quadratic fits' three large-*$n$ *estimates now span only* 0.602804 *to* 0.602822 *(versus* 0.602700 *to* 0.602782 *before this extension), and the two cubic estimates agree to five decimal places (*0.602824, 0.602825*). Second, and more tellingly, the* linear *model's own restricted-range estimate keeps climbing toward the quadratic/cubic value as the range is pushed to larger* $n$ — $0.6015 \to 0.6021 \to 0.6024$ *(proposition 25.18,* $n \leq 50$*)* $\to 0.602601 \to 0.602704 \to 0.602727$ *(proposition 25.42, extending to* $n = 100$*) — the same monotone, one-directional drift section 25.7.5 found in the log-log slope, now visible in a different fit entirely. Third, the gap between the linear model's best estimate (*0.602727*,* $n = 60$*-to-*100*) and the quadratic/cubic models' best estimates (*0.602822–0.602825*) has shrunk to about* $1 \times 10^{-4}$*, an order of magnitude smaller than section 25.7.20's* $6 \times 10^{-4}$ *gap before this extension. This is exactly the behavior expected if* $\kappa(n) = \kappa_\infty + c/n + d/n^2 + \cdots$ *genuinely converges and the linear model is simply omitting a real, shrinking-but-nonzero* $d/n^2$ *term: more data at larger* $n$ *makes the omission matter less, without requiring the underlying model to be revised. It is not a proof of convergence — no finite amount of numerical data can be one — but it is a second independent piece of evidence for it, on top of section 25.7.5's monotone log-log-slope narrowing, from a construction (fitting* $\kappa(n)$ *rather than* $\log \delta_{\text{pent}}(n)$*) that does not share that construction's own sources of noise.*

**Proposition 25.43** (The closed-form screen repeats at tightened precision, still empty)**.** *Rerunning proposition 25.37's fifteen-candidate screen against the tightened estimate* $\kappa_\infty \approx 0.602823$ *(midpoint of proposition 25.42's best quadratic and cubic fits) finds the same result: every candidate misses, by margins essentially unchanged from the earlier, looser screen (the closest,* $1/\varphi$*, still misses by* $1.5 \times 10^{-2}$*), and an extended rational search to* $q \leq 50$ *finds no match either.*

**Epistemic Notice**

Section 25.7.27 narrows, but does not close, section 25.7.20's honest-bound caveat: the best current estimate is $\kappa_\infty \approx 0.60282$, stable to five significant figures across the quadratic and cubic fits' large-$n$ ranges, with the once-troubling linear/quadratic disagreement now understood as a finite-range artifact rather than a sign of model failure. This is a numerical bound refined by adding data, not a proof; section 25.7.20's todobox — that identifying $\kappa_\infty$ in closed form requires a continuum-limit analysis of the relative curvature-density profile near an icosahedral vertex, which this chapter has not attempted — still stands, now against a numerically sharper target.

### 25.7.28 A practical ceiling: $n = 150, 200$, and $\kappa_\infty$ to five figures

Two further frequencies, $n = 150$ and $n = 200$ (800,000 Goldberg vertices at $n = 200$, several minutes of computation on the hardware this chapter was developed on), extend table 25.7 once more and mark the practical ceiling this chapter reaches: mesh-construction and star-ordering cost grow roughly as $n^2$, and beyond this range the computation becomes expensive enough that the diminishing returns on $\kappa_\infty$'s precision no longer justify the cost.

| $n$ | 150 | 200 |
|---|---|---|
| Pentagon-corner defect (°) | 0.0009742022 | 0.0005466226 |
| $\kappa(n)$ | 0.608876 | 0.607358 |

Table 25.8: Pentagon-corner defect and $\kappa(n)$, plane-projected convention, extending table 25.7 to the final two frequencies this chapter tests. Total defect sum remains 720° to within $7 \times 10^{-8}$ degrees at both.

**Proposition 25.44** (Refits reaching $n = 200$: the quadratic and cubic models converge to the same six figures)**.** *With $n = 150, 200$ included: quadratically, $n = 50$-to-200 (seven points) gives $\kappa_\infty = 0.602823$ with every residual under $6.6 \times 10^{-8}$, and $n = 100, 150, 200$ (exactly three points, an exactly-determined quadratic) reproduces $\kappa_\infty = 0.602823$ with residuals at machine precision ($1 \times 10^{-15}$); cubically, $n = 50$-to-200 gives 0.602824 (max residual $9.0 \times 10^{-10}$) and $n = 30$-to-200 gives 0.602824 (max residual $9.6 \times 10^{-9}$). The linear model, still included for comparison, continues its one-directional climb — 0.602769 over $n = 60$-to-200, 0.602793 over $n = 100, 150, 200$ — narrowing its gap to the quadratic/cubic value to under $3 \times 10^{-5}$, a further order of magnitude tighter than section 25.7.27's already-narrowed gap.*

**Key Result**

*[$\kappa_\infty = 0.60282$, stable to five significant figures across every model that includes a $1/n^2$ term, at every range tested from $n \geq 30$ up] Every quadratic or cubic fit run in this chapter, over any range with $n \geq 30$ and reaching at least $n = 100$, returns $\kappa_\infty$ in the narrow band 0.602819 to 0.602825 — a spread of $6 \times 10^{-6}$, three orders of magnitude tighter than section 25.7.20's original $n \leq 50$ bound. Repeating proposition 25.37's screen against this final estimate, with the candidate list extended to eighteen entries and the rational search extended to $q \leq 100$, finds the same result as every earlier pass: no candidate matches, the closest ($1/\varphi$) still missing by two orders of magnitude more than the tolerance. Taken together with section 25.7.27's finding that the once-large linear/quadratic disagreement shrinks systematically as more large-n data are added, this chapter's final position is: $\kappa_\infty \approx 0.60282$, is very likely a genuine limit (the convergence evidence, while not a proof, is now consistent across three independent model orders and five decades of range in n), and has no closed form among the eighteen icosahedrally-motivated candidates and hundred-odd small rationals tried here.*

**Epistemic Notice**

This closes the numerical side of the $\kappa_\infty$ investigation this chapter opened in section 25.7.7. What remains open — unchanged in kind from section 25.7.20's todobox, only sharpened in target precision — is a proof that $\lim_{n\to\infty} \kappa(n)$ exists, and, independently, an analytic identification of its value via the continuum-limit curvature-density argument sketched there. Both are genuinely open; neither is attempted here. Section 26.10 should record the asymptotic-rate item in its final form for this chapter: the pentagon-corner exponent is $-2$ to overwhelming numerical confidence (via section 25.7.7's exact mean and section 25.7.28's tightly bounded, closed-form-resistant ratio $\kappa_\infty \approx 0.60282$), with the existence of $\lim \kappa(n)$ and the value's analytic origin the two items this chapter leaves for future work.

### 25.7.29 Past the practical ceiling: a localized, $O(1)$-per-$n$ reconstruction reaches $n = 200{,}000$

Section 25.7.28's "practical ceiling" was a property of that section's own construction, not of the question itself: `goldberg_defect.py`'s cost grows as $n^2$ because it builds the *entire* Goldberg polyhedron — all $20n^2$ faces, with Python-dictionary vertex deduplication and a graph-traversal star-ordering at every one of them — before reading off the sixty pentagon-corner defects definition 25.17's $\kappa(n)$ actually needs. This subsection observes that each of those sixty defects depends only on a fixed-size neighborhood of the icosahedral vertex it sits beside: the corner triangle itself (definition 25.62's $T_k$, in the notation of section 25.7.49) together with one further ring of triangles, spanning at most the two icosahedral faces sharing an edge through that vertex. That neighborhood's combinatorial shape does not depend on $n$ at all — only the barycentric weights $(i, j, k)/n$ placing its six points do — so it can be built directly, at whatever frequency $n$, in time independent of $n$, rather than assembled as an afterthought of a global $O(n^2)$ mesh. A new, independent script, `scripts/kappa_localized.py`, implements this: for each of the twelve original icosahedral vertices it builds only the five local six-point/four-triangle patches meeting there (the near-corner rows $i = n, n-1, n-2$ of the same barycentric grid $p(i, j) = (iA + jB + kC)/n$, $k = n-i-j$, that `subdivide()` already uses, rather than the full $0 \le i \le n$ grid), stitches the shared edge-points across adjacent patches by coordinate matching exactly as `subdivide()` already does, and then reuses section 25.7.2's own `order_star` and `plane_projected_angle` routines unchanged on this much smaller mesh.

**Proposition 25.45** (The localized reconstruction reproduces the global one exactly)**.** *At $n = 8, 10, 20$ — three of table 25.2's already-published, independently-verified data points — the localized script's pentagon-corner mean defect agrees with `goldberg_defect.py`'s global computation to within $4.4 \times 10^{-7}$ degrees at every frequency tested (spread across the sixty corners at each frequency: at most $1.5 \times 10^{-13}$ degrees, consistent with exact icosahedral symmetry to machine precision). This is a genuine cross-check between two independently written scripts sharing no code beyond the plane-projection primitive itself, not a restatement of one construction in different variables.*

Having established the two constructions agree, this subsection uses the localized one — now cheap enough to run at any $n$ a laptop can represent in floating point — to push far past $n = 200$.

**Epistemic Notice**

[title=A precision bug found and fixed along the way] The first attempt to push proposition 25.45's script to large $n$ produced visibly wrong results above roughly $n = 300$: $\kappa(n)$ stopped decreasing smoothly and began to wander. The cause was traced to `angle_at` and `plane_projected_angle`'s use of $\arccos(\hat{e}_1 \cdot \hat{e}_2)$ to recover an angle from two unit vectors — a formula whose derivative diverges at $\cos\theta = 1$, so that float64's $\sim 1 \times 10^{-16}$ rounding in the dot product is amplified into an angle error of order $\sqrt{1 \times 10^{-16}} \approx 1 \times 10^{-8}$ radians, *independent of how small the true angle is.* Once $n$ is large enough that the true angle itself falls below that floor — which happens well before $n = 300$, since the pentagon-corner defect is already down to $2 \times 10^{-3}$ degrees at $n = 100$ (table 25.7) — arccos returns pure noise, occasionally rounding all the way to exactly zero. Replacing it with the standard numerically stable alternative, $\theta = \text{atan2}(\|\hat{e}_1 \times \hat{e}_2\|,\ \hat{e}_1 \cdot \hat{e}_2)$, which has no such singularity at $\theta = 0$, restores clean, monotone behavior out to much larger $n$ (section 25.7.29 below). This bug does not affect any number already published in this chapter: every value through table 25.8's $n = 200$ was computed by the global script, whose defects at that range are large enough (down only to $5.5 \times 10^{-4}$ degrees) that `arccos`'s noise floor, while present, stays several orders of magnitude below the reported precision. It became visible only because this subsection's whole purpose is to push the defect many further orders of magnitude smaller.

**Key Result**

*[$\kappa(n)$, extended smoothly from $n = 30$ to $n = 200{,}000$ — a thousand-fold increase in frequency over section 25.7.28's ceiling] With the fix in place, $\kappa(n)$ decreases smoothly and monotonically across every frequency tested from $n = 30$ to $n = 200{,}000$ (selected values: $\kappa(1000) = 0.6037282$, $\kappa(10{,}000) = 0.6029139$, $\kappa(100{,}000) = 0.6028298$, $\kappa(200{,}000) = 0.6028240$), with the sixty-corner spread at every frequency remaining at machine precision ($\leq 1.5 \times 10^{-13}$ degrees), exactly as icosahedral symmetry requires. Fits of $\kappa(n)$ against $1/n$ of orders one through four, over every large-$n$ range tried from $n \geq 100$ to $n \geq 10{,}000$, all return $\kappa_\infty$ in the narrow band 0.602821 to 0.602823, with residuals bottoming out at $\sim 2 \times 10^{-6}$ to $3 \times 10^{-6}$ regardless of model order or range — a residual floor that does not shrink with more data or a higher-order fit, unlike proposition 25.44's earlier fits, and is accordingly read as a further, smaller manifestation of float64 rounding rather than as unmodeled structure in $\kappa(n)$ itself (see section 25.7.29 below for the point at which this same source of error stops being a minor floor and starts corrupting the data outright).*

**Epistemic Notice**

[title=A second, different breakdown, found and left unfixed] Pushing further — $n = 300{,}000$ and above — $\kappa(n)$ stops decreasing monotonically at all: $\kappa(300{,}000) = 0.6027851$, $\kappa(350{,}000) = 0.6027938$ (larger than the 300,000 value), $\kappa(500{,}000) = 0.6026857$, $\kappa(1{,}000{,}000) = 0.6024442$, a visible and accelerating departure from the smooth trend section 25.7.29 reports below $n = 200{,}000$. This is a different failure from the arccos bug just fixed: the atan2-based angle formula has no small-angle singularity, but the underlying *points* themselves become the problem at extreme $n$ — the near-corner grid points used here sit at barycentric weight $(n-2)/n$ on the apex, a value that is 1 to float64 precision once $n$ exceeds roughly $2 \times 10^5$ to $3 \times 10^5$, degrading the geometric distinctness of the six points each local patch depends on. No attempt is made here to fix this second issue (it would require arbitrary-precision arithmetic, e.g. `mpmath`, applied to the point construction itself, not merely to the angle formula) — it is recorded honestly as the new ceiling this subsection actually reaches, at $n = 200{,}000$, a thousand-fold improvement on section 25.7.28's $n = 200$ but not the unlimited extension the $O(1)$-per-$n$ construction might otherwise seem to promise.

**Proposition 25.46** (The closed-form screen, repeated once more at the tightened estimate)**.** *Repeating proposition 25.43's screen against $\kappa_\infty \approx 0.602822$ (the central value of section 25.7.29's stable band), with tolerance $5 \times 10^{-6}$ matched to that band's own width and the rational search extended to $q \leq 200$: the same result as every earlier pass, no candidate*

*among the eighteen icosahedrally-motivated closed forms already tried, and no rational $p/q$ with $q \leq 200$, falls within tolerance.*

**Epistemic Notice**

Taken together, this subsection neither closes nor reopens anything section 25.7.28 left standing: it is an independent corroboration, from a differently-constructed script reaching three further orders of magnitude in $n$, of the same value, $\kappa_\infty \approx 0.60282$, and of the same conclusion, that no closed form among the candidates tried here matches it. What it adds is two things this volume did not have before: a concrete demonstration that the earlier "practical ceiling" was a property of one particular (global, $O(n^2)$) construction rather than of the underlying question, and a second, independently discovered numerical-precision failure mode (point degeneracy at extreme barycentric weight, distinct from `arccos`'s small-angle singularity) that would need to be addressed, by higher-precision arithmetic rather than a better angle formula, before $n$ could be pushed past roughly $2 \times 10^5$–$3 \times 10^5$ by any method built the way this one is. Section 26.10 should record this as a narrowing of *confidence* in $\kappa_\infty \approx 0.60282$ (now independently reproduced across two constructions and five and a half decades of range in $n$, from $n = 2$ to $n = 200{,}000$), not a narrowing of the value's own uncertainty band, which remains exactly where section 25.7.28 left it.

### 25.7.30 The named next step, taken: arbitrary-precision arithmetic past the point-degeneracy ceiling

Section 25.7.29's disclaimerbox left a concretely scoped, un-taken next step: the float64 localized construction's second failure mode — near-corner grid points becoming indistinguishable once their barycentric weight on the apex rounds to 1.0 in double precision, rather than any defect in the angle formula itself — would need arbitrary-precision arithmetic applied to the point construction, not merely a better angle formula, before $n$ could be pushed past roughly $2 \times 10^5$–$3 \times 10^5$. This subsection takes that step.

**Proposition 25.47** (An arbitrary-precision reimplementation, validated against every prior construction)**.** *`scripts/kappa_localized_mpmath.py` reproduces section 25.7.29's construction verbatim in its combinatorics — the same twelve corner patches, the same barycentric grid points $(n-1, 0)$, $(n-1, 1)$, $(n-2, 0)$, $(n-2, 1)$, $(n-2, 2)$, the same plane-projected angle convention — but replaces every float64 operation with `mpmath` arbitrary-precision arithmetic (50–80 decimal digits, depending on the frequency reached), including the best-fit plane normal, computed as the eigenvector of smallest eigenvalue of the $3 \times 3$ point-covariance matrix (`mpmath.eigsy`) rather than `numpy`'s SVD, since `mpmath` exposes no general SVD routine. At $n = 8, 10, 20$ — the same three frequencies proposition 25.45 already validated the float64 localized script against — the arbitrary-precision script reproduces the published values exactly, to every digit `mpmath` computes at 50 decimal digits of working precision, with the sixty-corner spread at machine-*arbitrary*-*precision floor ($1.2 \times 10^{-48}$ degrees, i.e. exactly zero at the precision used, as icosahedral symmetry requires).*

**Key Result**

*[Past the float64 ceiling by five orders of magnitude: $\kappa(n)$ computed cleanly from $n = 1 \times 10^5$ to $n = 1 \times 10^{10}$, with no trace of section 25.7.29's breakdown] With* 60–80 *decimal digits of working precision (chosen comfortably above what is needed even at the largest n reached, and increased from* 60 *to* 80 *digits for the runs at $n \geq 3 \times 10^8$ once the smaller working precision began to leave visibly fewer than the target number of stable digits), proposition 25.47's script computes $\kappa(n)$ smoothly and monotonically decreasing at every one of $n = 1 \times 10^5, 2 \times 1 \times 10^5, 3 \times 1 \times 10^5, 5 \times 1 \times 10^5, 1 \times 10^6, 3 \times 1 \times 10^6, 1 \times 10^7, 3 \times 1 \times 10^7, 1 \times 10^8, 3 \times 1 \times 10^8, 1 \times 10^9, 3 \times 1 \times 10^9, 1 \times 10^{10}$, with the sixty-corner spread at every single frequency at or indistinguishable from the working precision's own floor ($7 \times 10^{-59}$ to $1 \times 10^{-79}$ degrees, i.e. exact icosahedral symmetry, not approximate agreement). Selected values:*

$$\kappa(1 \times 10^5) = 0.602832543385, \quad \kappa(1 \times 10^7) = 0.602823592447,$$
$$\kappa(1 \times 10^9) = 0.602823502938, \quad \kappa(1 \times 10^{10}) = 0.602823502124.$$

*None of the qualitative breakdown section 25.7.29 reported at $n = 3 \times 10^5$ and above in float64 — non-monotonicity, then visible acceleration away from the trend — reappears anywhere in this range: the arbitrary-precision values continue section 25.7.29's float64 trend smoothly across the exact frequencies where the float64 construction broke down, confirming the breakdown's diagnosis (point degeneracy under fixed 64-bit precision, not a genuine feature of $\kappa(n)$ itself) directly rather than only by argument.*

**Proposition 25.48** (Five independently-fitted point sets agree to eighteen significant figures)**.** *Fitting $\kappa(n) = \kappa_\infty + c_1/n + c_2/n^2 + \cdots$ by an* exactly-determined *polynomial-in-$1/n$ system (as many terms as data points, solved exactly by* `mpmath.lu_solve` *rather than by ordinary least squares, so that each fit either reproduces every input point exactly or fails outright, with no averaging to obscure disagreement) over five different point sets drawn from section 25.7.30's largest-n values — a quadratic through $n = 1 \times 10^8, 3 \times 1 \times 10^8, 1 \times 10^9$; a second quadratic through $n = 3 \times 1 \times 10^8, 1 \times 10^9, 3 \times 1 \times 10^9$; a third through $n = 1 \times 10^9, 3 \times 1 \times 10^9, 1 \times 10^{10}$; a cubic through the four largest frequencies; and a second cubic shifted down by one point — returns*

$$\kappa_\infty = 0.602823502033767116\ldots$$

*identically across all five fits, agreeing to every one of the eighteen significant figures reported. A quartic fit through the five largest frequencies agrees with this value to the same eighteen figures as well.*

**Epistemic Notice**

The agreement proposition 25.48 reports is unusually strong evidence *for the numerical value*, precisely because the fits are exactly-determined rather than least-squares: with ordinary least-squares fits (as every earlier fit in this chapter used), residual noise in the input data can be partly absorbed into the fitted coefficients, so close agreement between different least-squares fits is only moderately informative about how correct any one of them is. An exactly-determined fit has no such slack — it either reproduces its input points exactly (which every fit here does, to the working precision) or it does not — so five different exactly-determined fits, built from disjoint or overlapping subsets of thirteen input values computed independently at each frequency, landing on the identical eighteen-digit value is a much stronger form of internal consistency than proposition 25.44 and section 25.7.28's earlier least-squares agreements provided. It is still not a proof that $\kappa(n) \to \kappa_\infty$: it is extremely strong evidence that *if* $\kappa(n)$ converges to a limit expressible (locally, near $n = \infty$) as a power series in $1/n$ — which every one of this chapter's fits, from proposition 25.18 onward, has assumed without proving — then that limit's value, to eighteen significant figures, is 0.602823502033767116.

**Proposition 25.49** (The closed-form screen, repeated once more against eighteen digits rather than five)**.** *Repeating proposition 25.46's screen at proposition 25.48's new precision: none of the eighteen previously-tried icosahedrally-motivated candidates survives at a tolerance of* $1 \times 10^{-15}$ *(the closest,* $1/\varphi \approx 0.618034$*, misses by* $1.5 \times 10^{-2}$*, unchanged in absolute terms from every earlier pass, but now* 13 *orders of magnitude outside tolerance rather than* 2*); an extended rational search to* $q \leq 300$ *finds nothing closer than* $85/141 \approx 0.602837\ldots$*, missing by* $1.3 \times 10^{-5}$*; and* `mpmath`*'s own* `identify` *routine — a PSLQ-based inverse symbolic calculator, given the target value and each of six candidate constant sets (*$\{\pi\}$*,* $\{\sqrt{5}\}$*,* $\{\varphi\}$*,* $\{\pi, \sqrt{5}\}$*,* $\{\pi, \varphi\}$*,* $\{\sqrt{5}, \varphi\}$*,* $\{\sqrt{3}, \sqrt{5}\}$*,* $\{\pi, \sqrt{3}\}$*,* $\{e, \varphi\}$*,* $\{\sqrt{2}, \sqrt{5}\}$*) — returns no candidate expression for any of them, a negative result from a general-purpose search tool rather than only from this chapter's own hand-picked candidate list.*

**Epistemic Notice**

What this subsection changes, honestly stated: not whether $\kappa_\infty$ has a closed form (still unknown, and now screened against a considerably wider and more automated search than before, with the same negative outcome), and not whether $\lim_{n\to\infty} \kappa(n)$ exists (still unproved, resting as it has throughout this chapter on the assumption that a power-series-in-$1/n$ fit extrapolates correctly). What has genuinely changed is the numerical value's own precision: from section 25.7.28's five- (later six-) significant-figure band, 0.602819 to 0.602825, to proposition 25.48's eighteen-digit, internally-consistent value 0.602823502033767116. This is the concrete outcome section 25.7.29's named next step promised, delivered by exactly the tool it named (arbitrary-precision arithmetic, applied to the point construction rather than only the angle formula) and no other new idea. Section 26.10 should retire the "tighten the numerical bound" item this chapter has carried since section 25.7.20's original todobox: at eighteen digits of self-consistency across independent exactly-determined fits, further numerical tightening of $\kappa_\infty$ is not a productive next step for this chapter to pursue further; what remains open is exactly what section 25.7.20's todobox always said remained open — an analytic derivation of the value, via the continuum-limit disclination machinery (section 25.7.21) this chapter has still not obtained and, on reflection, correctly declined to fake — not a better decimal approximation of it.

### 25.7.31 The shape statistics, extended to $n = 100$: the plateau holds

Section 25.7.28's taxonomy-chapter update (section 26.10) noted that the pentagon-corner decay-rate question and section 25.7.6's shape-profile question are, on reflection, the same open question asked from two directions. This subsection makes that connection concrete rather than only asserted: the same meshes already built for section 25.7.27's $\kappa(n)$ extension (table 25.7) supply, at no further computation, the full defect distribution needed to extend table 25.5's three shape statistics to $n = 60, 70, 80, 100$.

| $n$ | 60 | 70 | 80 | 100 |
|---|---|---|---|---|
| $\mathrm{CV}(n)$ | 0.1331 | 0.1332 | 0.1332 | 0.1333 |
| $M(n)$ | 1.206 | 1.206 | 1.206 | 1.206 |
| $R(n)$ | 0.6179 | 0.6157 | 0.6141 | 0.6119 |

Table 25.9: Definition 25.14's three shape statistics, extending table 25.5 to $n = 60, 70, 80, 100$.

**Key Result**

*[Doubling the frequency range moves* CV *by three parts in a thousand and* $M$ *not at all to the precision reported] Between table 25.5's* $n = 50$ *row and table 25.9's* $n = 100$ *row — doubling the frequency, the same range over which* $\kappa(n)$ *itself moved from* 0.6211 *to* 0.6119*, a relative change of about* 1.5% *—* $\mathrm{CV}(n)$ *moves from* 0.1330 *to* 0.1333 *(a* 0.2% *change, smaller than the* 1.6% *change section 25.7.6 already reported between* $n = 20$ *and* $n = 50$*) and* $M(n)$ *does not move at all to three decimal places (*1.206 *at every one of* $n = 50, 60, 70, 80, 100$*). $R(n)$ continues its steady decline (*$0.6208 \to 0.6119$*), consistent with section 25.7.6's reading of $R(n)$ as the one statistic still visibly in motion, tracking $\kappa(n)$'s own slow approach to $\kappa_\infty \approx 0.60282$ (section 25.7.28) rather than a genuinely distinct quantity — indeed $R(n)$ and $\kappa(n)$ differ only in denominator (the computed mean of the non-pentagon-corner vertices versus the exact overall mean $36°/n^2$), and both are decreasing toward the same neighborhood for the same reason.* $\mathrm{CV}(n)$ *and* $M(n)$*, by contrast, have essentially stopped moving over a frequency range where the defect magnitude itself has fallen by a further factor of four: this is stronger evidence for section 25.7.6's plateau reading than the original* $n \le 50$ *data alone provided, obtained without any new construction beyond what section 25.7.27 had already built for a different purpose.*

**Todo — verify primary source**

Section 25.7.31's extension uses only the four new frequencies computed for section 25.7.27; it does not extend to $n = 150, 200$, where section 25.7.28's data would give two further points at essentially no marginal computational cost beyond what has already been spent, and could be added in a future pass without new construction. What section 25.7.31 still does not do is what section 25.7.6's own todobox already flagged as missing: a proof that $\mathrm{CV}(n)$ and $M(n)$ converge, as opposed to continuing to drift at a rate too slow for even a further doubling to $n = 100$ to detect, and a closed-form candidate for the limiting profile $h(v)$ itself. The latter is, on the reading section 25.7.20's todobox already gave, the same unsolved continuum-limit problem $\kappa_\infty$'s closed form would require — $\kappa_\infty$ is $h$ evaluated (in the appropriate limiting sense) at an icosahedral-vertex direction, divided by $h$'s spherical average, and a derivation of one would very likely yield the other.

### 25.7.32 The shape statistics, extended once more: to $n = 120, 150, 200$

The previous todobox's extension is now carried out. Unlike section 25.7.31's reuse of meshes already built for section 25.7.27's pentagon-corner-only computation, table 25.8's $n = 150, 200$ pass computed only the sixty pentagon-corner defects, not the full vertex distribution definition 25.14's three statistics require, so this subsection's numbers involved constructing the full $n = 120, 150, 200$ meshes and their complete defect distributions from scratch (respectively 288,000, 450,000, and 800,000 Goldberg vertices), rather than reusing anything already computed; $n = 120$ is included as an intermediate point between table 25.9's $n = 100$ and the next frequency this chapter has previously tested at all ($n = 150$).

| $n$ | 120 | 150 | 200 |
|---|---|---|---|
| $\mathrm{CV}(n)$ | 0.1333 | 0.1334 | 0.1334 |
| $M(n)$ | 1.206 | 1.206 | 1.207 |
| $R(n)$ | 0.6103 | 0.6088 | 0.6073 |

Table 25.10: Definition 25.14's three shape statistics, extending table 25.9 to $n = 120, 150, 200$.

**Key Result**

*[Quadrupling the frequency past table 25.9's $n = 100$ row moves* CV *and $M$ by less than a part in a thousand; $R$ continues its same slow decline] Between table 25.9's $n = 100$ row (*$\mathrm{CV} = 0.1333$, $M = 1.206$, $R = 0.6119$*) and table 25.10's $n = 200$ row — doubling the frequency again, over the same range table 25.8 shows the pentagon-corner defect itself falling from* $0.001895°$ *to* $0.0005466°$*, a further* 3.5*-fold drop —* $\mathrm{CV}(n)$ *moves from* 0.1333 *to* 0.1334 *at the precision reported (a change smaller than section 25.7.31's already-small $n = 50 \to 100$ move) and $M(n)$ moves from* 1.206 *to* 1.207*, its first visible change (to three decimal places) anywhere in tables 25.5, 25.9 and 25.10's combined eighteen-frequency record — too small to contradict section 25.7.6's plateau reading, but worth recording honestly as the one point in this whole extension where a statistic previously reported as exactly frozen (section 25.7.31's "*1.206 *at every one of $n = 50, 60, 70, 80, 100$") is not frozen at the next digit. $R(n)$, by contrast, keeps moving in the same direction and at a visibly similar relative rate as before (*$0.6119 \to 0.6103 \to 0.6088 \to 0.6073$ *at $n = 100, 120, 150, 200$), continuing to track $\kappa(n)$'s own slow decline toward $\kappa_\infty \approx 0.60282$ (section 25.7.28) exactly as section 25.7.31 already read it — at $n = 200$, $R(n) = 0.6073$ sits* 0.0045 *above $\kappa_\infty$'s own value, a gap consistent with $R(n)$ and $\kappa(n)$ approaching the same limit from the same side at a comparable rate, though this subsection does not compute $\kappa(n)$'s own $n = 120$ value to check the two ratios' gap directly at every frequency tested here.*

**Todo — verify primary source**

Section 25.7.32's $M(n)$ finding — a first visible move at $n = 200$, after five frequencies ($n = 50$ through 100) at an apparently exact 1.206 — narrows section 25.7.6's plateau reading rather than overturning it: $M(n)$ is evidently still drifting, just at a rate slow enough that table 25.9's four-frequency window could not detect it and table 25.10's three-frequency window barely can. This is a small but genuine correction to how firmly section 25.7.31 stated the $M(n)$ plateau, and is recorded here rather than smoothed over. What remains exactly as open as section 25.7.31's own todobox left it: no proof that any of the three statistics converges, as opposed to continuing an ever-slower drift indefinitely, and no closed-form candidate for the limiting profile $h(v)$.

### 25.7.33 A first attempt at a convergence rate for $\mathrm{CV}(n)$ and $M(n)$, and why it does not yet succeed

Section 25.7.32's todobox leaves the convergence question entirely qualitative ("essentially stopped moving," "still drifting"). This subsection makes a first attempt to quantify it directly, by fitting a power-law model $\mathrm{stat}(n) = L_\infty - c/n^p$ to the twenty already- computed frequencies of tables 25.5, 25.9 and 25.10 combined, and reports both what the fit finds and why the finding should not yet be trusted.

**Proposition 25.50** (A tail fit ($n \geq 16$) is numerically well-behaved and suggestive; restricting further ($n \geq 50$) is not)**.** *`scripts/shape_stats_drift_fit.py` fits the model above by nonlinear least squares. Over $n \geq 16$ (fifteen points), the fit converges cleanly for both statistics:* $\mathrm{CV}_\infty = 0.133423 \pm 0.000018$ *with rate exponent $p = 1.883 \pm 0.040$, and $M_\infty = 1.206435 \pm 0.000255$ with $p = 1.790 \pm 0.374$ – both exponents intriguingly close to, but not pinned at, the pentagon-corner defect's own exactly-derived $p = 2$ (section 25.7.7), which would be the natural guess if a single $n^{-2}$ scale $\kappa(n)$ governs the whole field's approach to its limiting shape exactly as section 25.7.6's $\delta_n(v) \approx \kappa(n)h(v)$ picture suggests. Restricting to the more reliable large-n tail alone ($n \geq 50$, eight points) – the range where the model should matter most, and where section 25.7.7's own analogous restriction to the pentagon-corner defect* improved *the fit by an order of magnitude – does the opposite here: the CV fit drifts to a different, less well-constrained exponent ($p = 1.273$, no error bar reported by the solver), and the $M(n)$ fit fails outright, returning a numerically nonsensical intercept ($L_\infty \approx 3.33$, $p \approx 0$) driven by fitting three free parameters to a handful of points whose reported values change by only one unit in the last decimal place across the entire range.*

**Epistemic Notice**

Proposition 25.50's instability is not a failure of the fitting procedure; it is the expected consequence of fitting a three-parameter model to data reported at four decimal digits when the quantity being fit – the drift itself – is of order $1 \times 10^{-4}$ over the $n \geq 50$ range, i.e. at the rounding floor of the published table values themselves. Tables 25.5, 25.9 and 25.10's entries were never computed or stored, in any script this project retains, to more than that precision, so this is a genuine data limitation, not merely a display one: refitting with more decimal digits is not possible without recomputing the underlying meshes from scratch. An attempt to do exactly that – rebuilding the full defect distribution at $n = 300$, one frequency past section 25.7.32's largest, at full double precision, to obtain at least one high-precision anchor point for a more trustworthy fit – was made in the same pass as this subsection but did not complete within the computational budget available (the $n \geq 100$ full-distribution construction, using this chapter's existing non-vectorized per-vertex angle routine, scales too poorly to reach $n = 300$'s 1.8 million Goldberg vertices in reasonable time). This is recorded as an incomplete, not abandoned, next step: a vectorized rewrite of `full_defect_distribution()`, or simply budgeting a long-running batch pass rather than an interactive one, is a concretely scoped task for a future pass, not a research gap requiring new ideas.

**Todo — verify primary source**

This subsection's honest conclusion is negative: the convergence-rate question section 25.7.32's todobox left open is not answered by fitting the existing published data more cleverly, because that data was never recorded to the precision the question requires. The $n \geq 16$ fit's suggestive closeness to $p = 2$ is worth keeping in mind as a motivating guess for a future, higher-precision pass, but proposition 25.50's own $n \geq 50$ instability check shows plainly that this guess is not yet supported by evidence stable under a reasonable robustness check, and should not be reported as a finding on its own. The genuinely open items are therefore unchanged from section 25.7.32's own todobox in kind, but now sharpened in what would resolve them: (i) a double- precision recomputation of the full defect distribution at one or more frequencies past $n = 200$ (the $n = 300$ attempt above, completed rather than abandoned); (ii) refitting proposition 25.50's model against that higher-precision data; and, unchanged, (iii) any closed-form candidate for the limiting profile $h(v)$ itself, which this subsection, like every one before it in this section, does not attempt.

### 25.7.34 The postponed recomputation, completed: a vectorized rewrite reaches $n = 400$, and the convergence rate is no longer a guess

The item this section's closing todobox left as "concretely scoped" was completed in a later pass on the same topic, not a future one: `scripts/shape_stats_fast.py` rewrites `full_defect_distribution()` to build each primal vertex's cyclic star and plane-fit *once* rather than once per incident face (proposition 25.50's disclaimerbox already diagnosed the $O(V_{\text{primal}} \cdot \deg^3)$ redundancy this removes), cross-checks against tables 25.5, 25.9 and 25.10 at $n = 50, 100, 200$ (agreement to every digit those tables report), and then reaches $n = 250$, $n = 300$, and $n = 400$ (3,200,000 Goldberg vertices) in under three minutes each – fast enough that the true bottleneck was never the mesh size, only the earlier routine's unnecessary repetition.

**Key Result**

*[CV(n) and M(n) converge as $n^{-2}$; R(n) converges as $n^{-1}$ and its limit is $\kappa_\infty^{\mathrm{SSB}}$, because away from the sixty pentagon corners the mean defect converges to the naive flat baseline almost exactly] Fitting* $\mathrm{stat}(n) = L_\infty - c/n^p$ *to the six self-consistently recomputed points* $n = 50, 100, 200, 250, 300, 400$ *gives, for* $\mathrm{CV}(n)$, $p = 1.974 \pm 0.004$ *and* $L_\infty = 0.133418$, *with a maximum fit residual of* $3.5 \times 10^{-7}$ *– three orders of magnitude tighter than proposition 25.50's* $n \geq 16$ *fit to the rounded published table, and stable under dropping* $n = 50$ *(*$p = 1.972$*). For* $M(n)$ *the fit gives* $p = 2.040 \pm 0.088$, $L_\infty = 1.206563$, *noisier (residual* $8.3 \times 10^{-6}$, *reflecting* $M(n)$*'s known dependence on a single extremal vertex rather than a distributional average) but consistent with the same exponent. Both are compatible with the guess proposition 25.50 flagged as unconfirmed –* $p = 2$, *the pentagon-corner defect's own exactly derived exponent (section 25.7.7) – and this time the guess survives the robustness check that broke it before.*
*R(n) is qualitatively different: its fit gives* $p = 0.997 \pm 0.0002$, *i.e.* $R(n)$ *converges an order of magnitude* more slowly, *as* $n^{-1}$ *rather than* $n^{-2}$, *with* $L_\infty = 0.602815$. *This is not a new independent constant: it agrees with the independently established* $\kappa_\infty^{\mathrm{SSB}} \approx 0.602711$ *(section 25.7.8) to four significant figures (relative difference* 0.017%, *consistent with the power-law model's own form-misspecification error, since the true asymptotic form is known elsewhere in this chapter to mix* $1/n$ *and* $1/n^2$ *terms rather than being a single power). The reason is direct:* $R(n) = \overline{\delta}_{\mathrm{pent}}(n)/\overline{\delta}_{\mathrm{other}}(n)$ *by definition (section 25.7.32), and a direct computation of the two means separately (not just their ratio) at* $n = 200, 300, 400$ *shows* $\overline{\delta}_{\mathrm{other}}(n)/(36/n^2) = 1.000029,\ 1.000013,\ 1.000007$ *respectively – converging to exactly* 1, *and roughly twice as fast in the exponent as* $R(n)$ *itself decays (consistent with an* $O(1/n^2)$ *correction, matching* $CV(n)$ *and* $M(n)$*'s own exponent). In other words: away from the sixty pentagon-corner vertices, the mean defect converges to the naive flat/uniform baseline* $36/n^2$ *almost immediately, with essentially no scheme-specific curvature-concentration correction surviving into the limit;* only *the pentagon-corner subset carries a nontrivial asymptotic ratio, and that ratio is exactly* $\kappa_\infty^{\mathrm{SSB}}$ *by definition.* $R(n)$*'s slower,* $O(1/n)$ *approach to its limit is therefore not an independent phenomenon requiring its own explanation: it is inherited entirely from* $\kappa_{\mathrm{pent}}(n)$*'s own already-known* $O(1/n)$*-dominated convergence (section 25.7.8), diluted by an essentially converged denominator.*

**Epistemic Notice**

This clarifies, but does not itself derive, section 25.7.32's $h(v)$: it establishes that $h(v)$ is asymptotically supported only on the 60 pentagon-corner vertices (every other vertex's contribution to the limiting shape distribution vanishes at the same $O(1/n^2)$ rate $\mathrm{CV}(n)$ and $M(n)$ themselves converge at, i.e. contributes nothing to the *limiting* distribution's shape, only to how fast that limit is approached) – a genuine narrowing of what $h(v)$ could be, since it rules out any limiting profile with non-trivial support away from the pentagon corners, but it says nothing about $h(v)$'s value *at* those sixty points beyond the already-known $\kappa_\infty^{\mathrm{SSB}}$ itself, and in particular gives no closed form.

### 25.7.35 The direct argument, found: an exact closed form for $\kappa_{\mathrm{other}}(n)$, not merely its asymptotics

The question the previous subsection's todobox left open — is there a direct argument for the non-pentagon-corner defect's approach to the flat baseline, rather than reading it off numerically after the fact — turns out to already be answered by two results this chapter proved much earlier for entirely different purposes, combined by nothing more than bookkeeping.

**Proposition 25.51** (Exact closed form for $\kappa_{\mathrm{other}}(n)$)**.** *For every n in definition 25.5's class-I construction,*

$$\kappa_{\mathrm{other}}(n) \;=\; \frac{n^2 - 3\,\kappa_{\mathrm{pent}}(n)}{n^2 - 3},$$

*exactly, where* $\kappa_{\mathrm{pent}}(n)$ *is the common value (shared by all sixty pentagon-corner vertices, table 25.1's "uniform" column and confirmed here to* $5 \times 10^{-14}$ *degrees at* $n = 100$*) of the*

*pentagon-corner defect divided by $36°/n^2$, i.e. exactly the $\kappa(n)$ already tabulated throughout section 25.7.8.*

*Proof.* Section 25.7.7 gives $\sum_v \delta(v) = 720°$ and $V(n) = 20n^2$ exactly, for every $n$, under the plane-projected convention this section uses throughout (proposition 25.11, extended to arbitrary $n$ by the $n$-independent holonomy argument of proposition 25.7's proof). Split the sum into its pentagon-corner and non-pentagon-corner parts: with exactly sixty pentagon-corner vertices at every $n$ (five per icosahedral vertex, twelve icosahedral vertices, no double-counting for $n$ large enough that adjacent pentagons do not share a Goldberg vertex, which holds throughout the range tables 25.5, 25.9 and 25.10 tests) all sharing one common value $\delta_{\text{pent}}(n)$ by symmetry,

$$720° = 60\,\delta_{\text{pent}}(n) + (20n^2 - 60)\,\overline{\delta}_{\text{other}}(n).$$

Substitute $\delta_{\text{pent}}(n) = \kappa_{\text{pent}}(n) \cdot 36°/n^2$ and solve for $\overline{\delta}_{\text{other}}(n)$:

$$\overline{\delta}_{\text{other}}(n) = \frac{720° - 2160°\,\kappa_{\text{pent}}(n)/n^2}{20n^2 - 60} = \frac{36°}{n^2} \cdot \frac{n^2 - 3\kappa_{\text{pent}}(n)}{n^2 - 3}.$$

Dividing by $36°/n^2$ gives the claimed formula for $\kappa_{\text{other}}(n) := \overline{\delta}_{\text{other}}(n)/(36°/n^2)$. Every step is exact algebra on two already-proven exact identities and one symmetry fact; no new geometric computation is required. □

**Key Result**

*[This is the direct argument the previous todobox asked for, and it fixes the exponent exactly, not just numerically] Proposition 25.51's formula, evaluated with `scripts/shape_stats_fast.py`'s independently measured $\kappa_{\text{pent}}(200) = 0.607358$, $\kappa_{\text{pent}}(300) = 0.605844$, $\kappa_{\text{pent}}(400) = 0.605087$, predicts $\kappa_{\text{other}}(200) = 1.0000294$, $\kappa_{\text{other}}(300) = 1.0000131$, $\kappa_{\text{other}}(400) = 1.0000074$ – matching the directly measured values 1.000029, 1.000013, 1.000007 (section 25.7.34) to every digit those measurements report. Because $\kappa_{\text{pent}}(n) \to \kappa_\infty^{\text{SSB}} \approx 0.602711$ (section 25.7.8), expanding the exact formula for large $n$ gives*

$$\kappa_{\text{other}}(n) = 1 + \frac{3\,(1 - \kappa_{\text{pent}}(n))}{n^2 - 3} = 1 + \frac{3(1 - \kappa_\infty^{\text{SSB}})}{n^2} + O(n^{-4}) \approx 1 + \frac{1.1919}{n^2},$$

*which is not merely consistent with section 25.7.34's numerically fitted exponent $p = 1.974$–$2.040$ but* proves *the exponent is exactly 2, with an exact leading coefficient expressed in closed form in terms of $\kappa_\infty^{\text{SSB}}$ itself – the one quantity in this whole chapter that still has no closed form of its own. This answers the previous todobox's question affirmatively and completely: the reason $\kappa_{\text{other}}(n) \to 1$ as $O(1/n^2)$ is not a separate geometric fact requiring its own four-point-style local argument at all, but a direct, algebraic consequence of curvature conservation ($\sum_v \delta(v) = 4\pi$, already proven) once the pentagon-corner contribution's own known size is subtracted off.*

**Epistemic Notice**

Proposition 25.51 says nothing new about $\kappa_{\text{pent}}(n)$ or its limit $\kappa_\infty^{\text{SSB}}$ itself – it is precisely because those remain uncomputed in closed form that section 25.7.35's expansion can state $\kappa_{\text{other}}(n)$'s leading coefficient only in terms of $\kappa_\infty^{\text{SSB}}$, not as a bare number. Nor does the argument extend, as given, to $M(n)$ or $\text{CV}(n)$ directly: both are nonlinear statistics of the full defect distribution (a maximum and a standard deviation, not a mean), and $\sum_v \delta(v) = 4\pi$ constrains only the distribution's first moment exactly. That $M(n)$ and $\text{CV}(n)$ empirically share the same $n^{-2}$ exponent (section 25.7.34) is therefore still only numerically supported, not proved by this subsection's argument, though the mechanism just established – a fixed total split between a shrinking, symmetric sixty-point set and an increasingly dominant, increasingly uniform remainder – is a plausible qualitative reason to expect the same rate to govern the higher moments too.

**Todo — verify primary source**

Proposition 25.51 closes the specific question the previous subsection's todobox posed. What it opens is the natural next one: section 25.7.34's $M(n)$ and $\mathrm{CV}(n)$ exponents are now the only piece of this section's numerical evidence still resting on a fit rather than a proof. A genuinely analogous argument for the *second moment* of the non-pentagon-corner defect distribution – not just its mean, which proposition 25.51 now handles exactly – would need a second exact conservation law beyond $\sum_v \delta(v) = 4\pi$ (a "$\sum_v \delta(v)^2$" identity has no known topological source the way the first moment does, and this chapter does not conjecture one). Absent that, $\mathrm{CV}(n)$ and $M(n)$'s exponent remains at the same evidential status section 25.7.34 left it: numerically indistinguishable from 2, not proved equal to it. Unchanged, as always: (iii) a closed-form value for $\kappa_\infty^{\mathrm{SSB}}$ itself.

### 25.7.36 The second moment: no new identity is needed, because one was already in hand

The preceding todobox asked for a "$\sum_v \delta(v)^2$" conservation law to put $\mathrm{CV}(n)$'s exponent on the same exact footing proposition 25.51 gave the mean. This subsection shows no new law is needed: the second moment $\sum_v \delta(v)^2$ is already fixed, exactly and at every finite $n$, by the two facts this chapter already has in hand – $N(n) = 20n^2$ (proposition 25.16) and $\sum_v \delta(v) = 4\pi$ (proposition 25.7) – together with $\mathrm{CV}(n)$'s own definition, purely algebraically and without computing a single new angle.

**Proposition 25.52** (The second-moment identity)**.** *For every $n \geq 2$, writing $\bar{\delta}(n) := \frac{1}{N(n)} \sum_v \delta(v)$ for the mean defect and $\mathrm{CV}(n) := \mathrm{std}(\delta)/\bar{\delta}(n)$ for the coefficient of variation over $\mathcal{P}_n$'s $N(n) = 20n^2$ vertices (both as in section 25.7.6), the following holds exactly, with no asymptotic qualifier:*

$$\sum_v \delta(v)^2 \;=\; N(n)\,\bar{\delta}(n)^2\big(1+\mathrm{CV}(n)^2\big) \;=\; \frac{(4\pi)^2}{20n^2}\big(1+\mathrm{CV}(n)^2\big) \;=\; \frac{4\pi^2}{5n^2}\big(1+\mathrm{CV}(n)^2\big).$$

*Proof.* For any finite set of $N$ numbers with mean $\bar{\delta}$ and standard deviation $\sigma$, $\sum_v \delta(v)^2 = N(\sigma^2+\bar{\delta}^2) = N\bar{\delta}^2(1+(\sigma/\bar{\delta})^2) = N\bar{\delta}^2(1+\mathrm{CV}^2)$ by the ordinary variance identity $\sigma^2 = \overline{\delta^2} - \bar{\delta}^2$; this holds for any finite population and is not specific to $\mathcal{P}_n$. Substituting $\bar{\delta}(n) = 4\pi/N(n)$, itself exact by proposition 25.7's numerator and proposition 25.16's denominator with no fitting or limit involved, and then $N(n) = 20n^2$, gives the stated closed form. □

**Key Result**

*[Verified to twelve significant figures, and the previous todobox's premise was false] Proposition 25.52's right-hand side, evaluated from each frequency's own measured $\mathrm{CV}(n)$, reproduces the independently measured $\sum_v \delta(v)^2$ to a relative error of $8.4 \times 10^{-14}$ at $n = 10$, growing slowly with the number of terms summed (as expected of floating-point accumulation, not of any discrepancy in the identity) to $3.2 \times 10^{-11}$ at $n = 200$ (script `second_moment_identity.py`) — agreement at the level of double-precision roundoff itself, not a numerical coincidence needing further explanation. Consequently, the previous todobox's stated premise – that a "$\sum_v \delta(v)^2$" identity "has no known topological source the way the first moment does" – was not quite the right way to put the gap: no* topological *source is needed, because an* algebraic *one, resting only on results section 25.7.7 already proved, suffices. What this does and does not deliver should be stated with the same care proposition 25.51's own keyresult used: it converts $\sum_v \delta(v)^2$'s $O(1/n^2)$ decay into an exact restatement of $\mathrm{CV}(n) \to \mathrm{CV}_\infty$, but it does not prove that limit exists, nor supply $\mathrm{CV}_\infty$'s value – exactly the two things section 25.7.34's fit and this chapter's whole $\kappa(n)$ cascade were already unable to supply, and still cannot. The second moment and $\mathrm{CV}(n)$'s convergence were never two separate open questions needing two separate conservation laws; proposition 25.52 makes visible that they were the same question stated twice.*

**Todo — verify primary source**

What proposition 25.52 does *not* do: it says nothing about *why* $\mathrm{CV}(n)$ should converge to a finite limit at all, let alone at rate $O(1/n^2)$ (section 25.7.34's own fitted exponent, still a fit and not a proof) – that remains exactly as open as $\kappa_\infty^{\mathrm{SSB}}$ itself, and for a related reason: both ask for the limiting *shape* of the defect distribution across an icosahedral vertex star, not merely its total. $M(n)$, unlike $\mathrm{CV}(n)$, involves $\max_v \delta(v)$ rather than a moment, so no analogous algebraic shortcut is available for it from $N(n)$ and $\sum_v \delta(v)$ alone; its own $O(1/n^2)$-consistent exponent (section 25.7.34) remains a bare numerical fit, with no identity of any kind behind it. Unchanged, as always: a closed-form value for $\kappa_\infty^{\mathrm{SSB}}$ itself.

### 25.7.37 Why $\kappa_\infty^{\mathrm{SSB}}$'s localization trick cannot be transplanted to $M(n)$: the maximizing vertex is not corner-local at all

The preceding todobox left $M(n)$'s exponent as "a bare numerical fit, with no identity of any kind behind it," and left open, without stating it as a concrete next step, whether the same trick that pushed $\kappa_\infty^{\mathrm{SSB}}$ from a practical ceiling of $n = 200$ to $n = 200{,}000$ (section 25.7.29's $O(1)$-per-$n$ local patch, exploiting the fact that a pentagon-corner defect depends only on a fixed-size neighborhood of its icosahedral vertex) could do the same for $M(n)$. That trick requires one structural fact to hold: that the Goldberg vertex actually attaining $\max_v \delta(v)$ sits at a fixed combinatorial position relative to the nearest of the twelve original icosahedral vertices, independent of $n$. `scripts/m_n_argmax_locus.py` tests this directly, by tagging every dual (Goldberg) vertex of section 25.7.34's own fast full-mesh construction with its ring-distance $r$ (in subdivision steps) to the nearest of the three icosahedral corners of the face it sits in, then reading off $r$ at the vertex actually achieving $\max_v \delta(v)$, at $n = 20, 50, 100, 150, 200, 300$.

**Proposition 25.53** ($M(n)$'s maximizing locus scales with $n$; it does not sit at fixed distance from a corner)**.** *At every tested $n$, the maximizing Goldberg vertex (or, for $n \geq 100$, one of a growing tied or near-tied set — 80 faces within $1 \times 10^{-6}$ of the maximum at $n = 100$, 2060 at $n = 200$) sits at ring-distance $r(n)$ from its nearest icosahedral corner with*

$$r(n) \in \{6,\ 16,\ 33,\ 49,\ 66,\ 99\} \text{ for } n \in \{20, 50, 100, 150, 200, 300\},$$

*giving $r(n)/n = 0.300,\ 0.320,\ 0.330,\ 0.327,\ 0.330,\ 0.330$ — monotonically approaching, and by $n = 100$ already agreeing to three decimal places with, $1/3$. Since $r/n = 1/3$ is exactly the location of a face's own barycentric centroid ($i = j = k = n/3$), this identifies the maximizing locus not as a fixed neighborhood of an icosahedral corner but as one that migrates, at a fixed* fraction *of the way across each face, toward that face's centroid as $n \to \infty$.*

**Epistemic Notice**

This rules out, rather than confirms, the natural next step: section 25.7.29's $O(1)$-per-$n$ construction is only possible because $\delta_{\text{pent}}(n)$ depends on a patch of *fixed* size around a corner, so that patch can be built directly at any $n$ without touching the other $O(n^2)$ faces. Proposition 25.53 shows $M(n)$ has no such patch: its maximizing vertex's distance from the nearest corner grows linearly in $n$, so any local construction reaching it would itself have to grow with $n$ — at which point it is simply the full mesh again, at whatever cost section 25.7.34's vectorized routine already achieves. This is a negative result, but an informative one: it explains, for the first time in this cascade, *why* $M(n)$'s exponent has remained a bare fit while $\kappa_\infty^{\text{SSB}}$'s own value became reproducible to five and a half decades of range in $n$ — the two quantities are not different only in how hard they have been to pin down numerically, but in what kind of geometric object governs each: $\kappa_\infty^{\text{SSB}}$ is a genuinely local property of the disclination at a single icosahedral vertex, while $M(n)$, on this evidence, is closer to a property of the flat-chord interpolation error *across an entire face*, maximized somewhere in its interior rather than at either type of vertex the chapter has so far singled out (pentagon-corner or first-ring hexagon-only). No claim is made here about $\mathrm{CV}(n)$, a distributional statistic over all $N(n)$ vertices rather than a single extremal one; proposition 25.53's argument does not bear on it one way or the other, and its own convergence remains exactly as open as section 25.7.36's todobox left it.

**Todo — verify primary source**

What this section changes: $M(n)$'s open item is no longer merely "the exponent is a bare fit"; it is now, more usefully, "the exponent is a bare fit *because* the extremal locus is a whole-face phenomenon at relative position $\approx 1/3$, not a corner-local one, and no analytic model of flat-chord interpolation error across the interior of a spherical triangle has been attempted in this volume." A genuine derivation of $M(n)$'s rate would need exactly that model — the deviation of the SSB chord grid from the true spherical geodesic grid as a function of barycentric position within a face, maximized over that position — rather than an extension of section 25.7.29's corner-local method, which this section shows cannot reach it even in principle. Unchanged, as always: whether $M(n)$ and $\mathrm{CV}(n)$ converge at all (as opposed to drifting indefinitely at an ever-slowing rate), and a closed-form value for $\kappa_\infty^{\text{SSB}}$ itself.

### 25.7.38 A closed-form candidate for $M_\infty$: the flat-chord area Jacobian, computed exactly

Section 25.7.37's negative result — that $M(n)$'s maximizing locus is a whole-face phenomenon at relative position $1/3$, not a corner-local one — is not only a dead end for localization; it identifies exactly which continuum computation would settle $M(n)$'s rate and limit, if the same "leading-order Gauss–Bonnet" area hypothesis that explained most (though not all) of $\kappa_\infty$'s departure from 1 (section 25.7.24) also governs $M(n)$. This subsection carries that computation out, and finds it tractable in closed form precisely because the relevant point — a face's own centroid — is fixed and symmetric, unlike $\kappa_\infty^{\text{SSB}}$'s corner, which needs the full disclination machinery this chapter has not obtained.

**Proposition 25.54** (The SSB area-distortion Jacobian, in closed form)**.** *Let $A, B, C \in S^2$ be one icosahedral face's three (unit) vertices, and let $p(\alpha,\beta,\gamma) = \alpha A + \beta B + \gamma C$ for $\alpha+\beta+\gamma = 1$, $\alpha,\beta,\gamma \geq 0$, parametrize SSB's flat chord grid on that face (section 25.1's barycentric construction, before the single final radial projection $P = p/|p|$ this scheme applies). Writing $r(\alpha,\beta,\gamma) := |p|$, $\hat{n}$ for the plane's unit normal, and $h := \hat{n}\cdot A$ (the plane's signed distance from the origin, the same for every point of the plane since $A, B, C$ all lie on it), the ratio of the sphere's area element (pulled back through $P$) to the flat grid's own uniform area element in $(\alpha,\beta)$ is*

$$J(\alpha,\beta,\gamma) \;=\; \frac{|h|}{r(\alpha,\beta,\gamma)^3}.$$

*Proof.* With $u := \partial p/\partial\alpha = A - C$ and $v := \partial p/\partial\beta = B - C$ (both constant, since $p$ is affine in $(\alpha, \beta)$), the flat area element is $|u \times v|\, d\alpha\, d\beta$, independent of $(\alpha, \beta)$. For the pushforward through $P = p/r$: $\partial P/\partial\alpha = r^{-1}\Pi(u)$ and $\partial P/\partial\beta = r^{-1}\Pi(v)$, where $\Pi(w) := w - (w \cdot P)P$ is orthogonal projection onto the tangent plane at $P$ (the standard derivative of unit-normalization). Since $u \times v$ is parallel to $\hat{n}$ with $u \times v = \pm|u \times v|\hat{n}$, and since $\Pi(u), \Pi(v)$ both lie in the plane orthogonal to $P$ so that $\Pi(u) \times \Pi(v)$ is parallel to $P$, its signed length is obtained by dotting with $P$: $P \cdot (\Pi(u) \times \Pi(v)) = P \cdot (u \times v)$ (the cross terms introduced by $\Pi$ each contain a scalar triple product with a repeated vector, and vanish). Substituting $P = p/r$ and $u \times v = \pm|u \times v|\hat{n}$ gives $P \cdot (u \times v) = \pm|u \times v|\,(\hat{n} \cdot p)/r = \pm|u \times v|\, h/r$ (using $\hat{n} \cdot p = h$ for every point of the plane, not only at $A$). Hence $|\partial P/\partial\alpha \times \partial P/\partial\beta| = r^{-2}|\Pi(u) \times \Pi(v)| = r^{-2} \cdot |h||u \times v|/r = |h|\,|u \times v|/r^3$, and dividing by the flat element $|u \times v|\, d\alpha\, d\beta$ gives the stated ratio. □

**Key Result**

*[$J$ is maximized exactly at the face centroid, and the maximum has a closed form] Since $A, B, C$ are unit vectors, $r = 1$ exactly at each corner ($\alpha = 1$ or $\beta = 1$ or $\gamma = 1$), and $r < 1$ strictly in the interior (a strictly convex combination of unit vectors has norm strictly less than 1 unless the vectors coincide). A regular icosahedral face carries a 3-fold rotational symmetry fixing only its own centroid $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ within the flat triangle, so $r$ has a critical point there; direct evaluation over the full simplex (a $101 \times 101$ grid, `scripts/m_n_area_hypothesis.py`) confirms this critical point is the* global minimum *of $r$ (equivalently, the global maximum of $J = |h|/r^3$) rather than a saddle, with the minimum value $r_{\min} \approx 0.794655$ matching $|(A + B + C)/3|$ to twelve digits. A short symmetry argument makes this exact rather than merely numerical: the same 3-fold rotation fixes $\hat{n}$ up to sign, so $\hat{n} \parallel (A + B + C)$, giving $h = \hat{n} \cdot A = |(A + B + C)/3| = r_{\min}$ exactly — the face's perpendicular foot from the origin and its barycentric centroid coincide, a consequence of the face's own symmetry rather than a general fact about triangles on a sphere. This collapses $J_{\max} = |h|/r_{\min}^3$ to $1/h^2$.*

*Using exact icosahedral coordinates $A = (0, 1, \varphi)/\sqrt{2 + \varphi}$ (and cyclic images) with $\varphi = (1 + \sqrt{5})/2$, symbolic evaluation (`sympy`) gives $h^2 = (5 + 2\sqrt{5})/15$ and the flat face area $\sqrt{3}(5 - \sqrt{5})/10$, hence, dividing $J_{\max} = 1/h^2$ by the mean of $J$ over the face (the sphere-to-flat area ratio, $(\pi/5)$/[flat face area], since 20 congruent faces tile the sphere's total area $4\pi$):*

$$M_\infty^{(\text{area})} := \frac{J_{\max}}{\overline{J}} = \frac{15(7\sqrt{3} - 3\sqrt{15})}{2\pi} = 1.2065670268765913\ldots$$

**Key Result**

*[Numerical test: the area hypothesis accounts for $M(n)$'s excess over 1 almost completely, and increasingly so as $n$ grows — unlike $\kappa_\infty$'s stable 4% residual]* `scripts/m_n_area_hypothesis.py` *computes both $M(n) = \max_v \delta(v)/\bar{\delta}(n)$ and the area-only analogue $M_{\text{area}}(n) := \max_v A(v)/\bar{A}(n)$ (dual-cell area, defined exactly as in section 25.7.24) at the same three frequencies:*

| $n$ | $M(n)$ | $M_{\text{area}}(n)$ | $(M_{\text{area}}(n)-1)/(M(n)-1)$ |
|---|---|---|---|
| 50 | 1.206039 | 1.205671 | 0.9982 |
| 100 | 1.206435 | 1.206343 | 0.9996 |
| 200 | 1.206534 | 1.206511 | 0.9999 |

*with the two argmax vertices — of defect and of area separately — landing at the identical ring position at every tested $n$ (section 25.7.37's $16, 33, 66$), and the vertex-by-vertex correlation between $\delta(v)$ and $A(v)$ across the* entire *distribution reaching* 0.999996 *at $n = 200$. The unexplained fraction (*0.18%, 0.04%, 0.01%*) shrinks by almost exactly a factor of* 4 *at each doubling of $n$ — an $O(1/n^2)$ decay, the same rate section 25.7.34 already established for $M(n)$'s own approach to its limit. This is the qualitative opposite of what sections 25.7.24 and 25.7.25 found for $\kappa_\infty^{\text{SSB}}$'s own pentagon-corner area hypothesis, where a genuine, non-vanishing $\sim 4\%$ residual survived every correction attempted (sections 25.7.25 and 25.7.26) and was traced to fivefold combinatorial valence specifically. Here, on present evidence, there is no analogous floor: the residual is not merely small but visibly heading to zero.*

**Proposition 25.55** (Independent confirmation: a direct fit of $M(n)$ itself lands within the fit's own noise of the closed form)**.** *Extending section 25.7.34's own five-point run by one further frequency, $n = 400$ (*`scripts/m_n_argmax_locus.py`*, $M(400) = 1.206559$), and refitting $M(n) = M_\infty - c/n^2$ by ordinary least squares on $n = 50, 100, 200, 300, 400$ gives $M_\infty^{(\text{fit})} = 1.2065637$ with maximum residual $8.2 \times 10^{-6}$ — a fit made independently of section 25.7.38's analytic derivation, using only the raw defect maxima. The two values,*

$$M_\infty^{(\text{fit})} = 1.2065637 \qquad \textit{vs.} \qquad M_\infty^{(\text{area})} = 1.2065670,$$

*differ by $3.3 \times 10^{-6}$ — smaller than the fit's own maximum residual, i.e. indistinguishable from it given the data this chapter has. This is the strongest evidence available in this volume that the closed form of section 25.7.38 is not merely a leading-order approximation to $M_\infty$ but its exact value.*

**Epistemic Notice**

What this section does and does not establish should be stated as carefully as section 25.7.24's own analogous claim for $\kappa_\infty$ was. Proposition 25.54 is a genuine proof, not a numerical finding: the Jacobian formula $J = |h|/r^3$ follows from the definition of central projection alone, with no appeal to $n \to \infty$ or to any property specific to the icosahedron beyond its face being planar. Section 25.7.38's identification of the maximizing point, and the resulting closed-form value $M_\infty^{(\text{area})}$, are likewise exact — *conditional on* the leading-order area hypothesis itself (that $\delta(v)/\bar{\delta}(n) \to A(v)/\bar{A}(n)$ as $n \to \infty$, vertex-by-vertex, not only in the aggregate) being not merely a good approximation but an *exact* statement at the specific point of section 25.7.37's migrating locus. That hypothesis is not proved here, in the sense of a rigorous bound on the gap between $\delta(v)$ and $K \cdot A(v)$ at a generic interior vertex as $n \to \infty$; section 25.7.38 and proposition 25.55 are strong numerical evidence for it at this specific point, not a substitute for such a bound, and the same style of evidence looked equally clean for $\kappa_\infty^{\text{SSB}}$ before section 25.7.25's later, more careful test found a genuine residual anyway. What distinguishes the present case, on the evidence actually in hand, is that the residual here is shrinking with $n$ rather than converging to a nonzero constant — but "shrinking over three data points" is exactly the kind of claim proposition 25.50 elsewhere in this chapter warned against trusting without a high-precision anchor, and this section's own anchor is $n = 400$, not $n = 200{,}000$.

**Todo — verify primary source**

A genuine proof that $M_\infty = 15(7\sqrt{3} - 3\sqrt{15})/(2\pi)$ exactly would need to rule out a residual at the face-centroid locus the way section 25.7.26 eventually ruled out (for the pentagon corner) a curved-versus-flat measurement artifact and section 25.7.25 ruled out (there, unsuccessfully) a smooth shape correction — i.e. a local expansion of $\delta(v) - K \cdot A(v)$ near $\alpha = \beta = \gamma = 1/3$ to the next order in $1/n$, checked against a much higher-precision computation than this chapter's $n = 400$ ceiling for the full-mesh routine. This is a concretely scoped next step, not a new research direction: section 25.7.29's $O(1)$-per-$n$ technique cannot reach it (section 25.7.37), but a second, differently-localized construction — exploiting that the centroid locus, unlike the corner locus, needs only a small neighborhood of *one interior grid point* rather than of an icosahedral vertex — has not yet been attempted and may make an equivalent push to very large $n$ possible by the same kind of trick, just anchored at a different point. Unchanged: whether $M(n)$ and $\mathrm{CV}(n)$ converge at all is still asserted only on the strength of section 25.7.34's fit, not proved; $\mathrm{CV}(n)$'s own limit has no analogous area-based candidate proposed *here* (section 25.7.39 below removes exactly this caveat, by observing that a variance needs no argmax); and $\kappa_\infty^{\mathrm{SSB}}$'s own closed form — where, unlike here, this same style of area argument is already known (section 25.7.24) to fall 4% short — remains exactly as open as ever.

### 25.7.39 A closed-form candidate for $\mathrm{CV}_\infty$ too: no argmax is needed for a variance

The preceding todobox's caveat was narrower than it needed to be. Section 25.7.37's obstruction to transplanting section 25.7.29's localization trick to $M(n)$ was specifically that $M(n)$ is an *argmax*, and the maximizing vertex has no fixed local neighborhood to localize around as $n \to \infty$. $\mathrm{CV}(n)$ is not an argmax: it is a ratio of two moments of the full defect distribution, and proposition 25.54's area-distortion Jacobian $J(\alpha, \beta, \gamma)$ already *is* the continuum-limit relative-defect profile the same leading-order area hypothesis proposes, at every point of a face, not only at its centroid. Under that hypothesis, as $n \to \infty$ the (rescaled) defect at a primal vertex converges to $J$ evaluated at that vertex's barycentric position, and since the SSB grid places one vertex per unit cell of $(\alpha, \beta)$-area, the empirical distribution of rescaled defects over one face converges to the distribution of $J$ under the *uniform* measure on the flat triangle — exactly the measure section 25.7.38's own denominator $\overline{J}$ already integrates over. Two moments of that same measure, not one optimization, are all $\mathrm{CV}_\infty^{(\mathrm{area})}$ needs.

**Key Result**

*[$\mathrm{CV}_\infty^{(\mathrm{area})}$, computed as the coefficient of variation of $J$ itself]* `scripts/cv_area_hypothesis.py` *computes $E[J]$ and $E[J^2]$ under the uniform measure on one icosahedral face's flat triangle in barycentric coordinates $(\alpha, \beta)$, first by adaptive double-precision quadrature (*`scipy.integrate.dblquad`*, absolute and relative tolerance $10^{-13}$) and independently by 50-digit* `mpmath` *quadrature, the two agreeing to every digit the double-precision run reports. $E[J]$ reproduces section 25.7.38's own denominator exactly, $E[J] = 2\pi/(5\sqrt{3} - \sqrt{15}) = 1.312477549714\ldots$, a consistency check the script performs automatically rather than merely asserts. $E[J^2] = 1.753260058940\ldots$ is new to this computation. The coefficient of variation of $J$ itself,*

$$\mathrm{CV}_\infty^{(\mathrm{area})} := \frac{\sqrt{E[J^2] - E[J]^2}}{E[J]} = 0.1334177461964\ldots$$

*matches section 25.7.34's independently fitted $\mathrm{CV}_\infty^{(\mathrm{fit})} = 0.133418$ to every digit that fit's own $3.5 \times 10^{-7}$ residual can support (relative difference $2 \times 10^{-4}\%$) — the same order of agreement proposition 25.55 found between $M_\infty^{(\mathrm{fit})}$ and $M_\infty^{(\mathrm{area})}$.*

**Epistemic Notice**

This is conditional on the same leading-order area hypothesis as section 25.7.38, and inherits every caveat section 25.7.38's own disclaimerbox already stated for $M_\infty^{(\text{area})}$: agreement this clean, at a single tested anchor, is exactly the kind of evidence that later turned out to hide a genuine few-percent residual for $\kappa_\infty^{\text{SSB}}$'s own corner-local version of the same hypothesis (section 25.7.25). What is different here, and worth stating plainly rather than eliding: $\text{CV}_\infty$'s comparison is against a single fitted number from six data points, not against a 200,000-frequency localized reconstruction of the kind sections 25.7.29 and 25.7.40 used to stress-test the corner and centroid cases respectively — no analogous $O(1)$-per-$n$ localized construction for the *whole-distribution* second moment has been attempted, and it is not obvious one exists, since $\text{CV}(n)$'s definition genuinely needs every vertex's defect, not a neighborhood of one. The agreement reported here is real and independently reproduced (scipy and mpmath, two unrelated quadrature implementations), but it rests on one comparison point, not a convergent sequence.

**Proposition 25.56** (A closed form for $\text{CV}_\infty^{(\text{area})}$ itself: screened and not found)**.** *Following proposition 25.37's own method for $\kappa_\infty$, `cv_area_hypothesis.py` tests $E[J^2]$, $\text{CV}_\infty^{(\text{area})}$, and $\left(\text{CV}_\infty^{(\text{area})}\right)^2$ against `mpmath.identify` over the same icosahedrally-motivated constant basis used throughout this chapter ($\pi$, $\sqrt{3}$, $\sqrt{5}$, $\sqrt{15}$, and $\pi^{-2}$-scaled combinations of each), and separately runs `mpmath.pslq` directly on $E[J^2]$ and $\left(\text{CV}_\infty^{(\text{area})}\right)^2$ against the six-element basis $\{1, \pi^2, \sqrt{3}\,\pi^2, \sqrt{5}\,\pi^2, \sqrt{15}\,\pi^2\}$ at 60 decimal digits. Every screen returns no relation, in contrast to $E[J]$ and $M_\infty^{(\text{area})}$, both of which are elementary combinations of exactly this basis. Unlike $\kappa_\infty^{\text{SSB}}$'s own screen (proposition 25.37), which tried eighteen icosahedrally-motivated candidates plus every low-denominator rational and is reported as a broad negative result, this is a narrower screen over one natural constant family; a genuinely closed form built from a different family (for instance, one involving the elliptic integrals a degree-three polynomial denominator under a two-dimensional domain integral would more naturally produce) is not ruled out by this test.*

**Todo — verify primary source**

$\text{CV}(n)$'s own convergence is therefore in the same position section 25.7.38 left $M(n)$ in: a strongly evidenced numerical candidate for the limit, conditional on the same unproved leading-order area hypothesis, with no closed form yet identified for the candidate itself. What remains open, specific to CV: (i) a higher-precision direct fit of $\text{CV}(n)$ itself, analogous to proposition 25.55's $n = 400$ anchor for $M_\infty$, to confirm the agreement is not a coincidence of section 25.7.34's particular six-point fit; (ii) whether $E[J^2]$ admits a closed form outside the constant family proposition 25.56 tested, given that $r^{-6}$ integrated over a flat triangle is exactly the kind of integral that produces elliptic integrals rather than the algebraic surds $E[J]$ and $M_\infty^{(\text{area})}$ reduced to; and (iii), as with $M_\infty$, a genuine bound on the gap between the area hypothesis and the true defect at a generic interior vertex, which this section — like section 25.7.38 before it — does not supply. $\kappa_\infty^{\text{SSB}}$'s own closed form remains, as ever, the more resistant case: the one place this chapter has found where the same area hypothesis is already known to leave a genuine, non-vanishing residual rather than an apparently exact match.

### 25.7.40 The centroid locus, localized after all: an $O(1)$-per-$n$ construction reaches $n = 10{,}000$, and a new, differently-caused ceiling

The preceding todobox named, as a concretely scoped next step, a second local construction anchored at the face-centroid locus rather than at an icosahedral corner. `scripts/m_centroid_localized.py` builds it: away from every icosahedral corner and every face edge, every primal (subdivided-icosahedron) vertex has the same combinatorial structure — a hexagonal star of six incident triangles — independent of $n$ except through the barycentric weights used to place its six points, exactly the observation section 25.7.29's

corner-local construction already exploited at a fivefold rather than sixfold star. A radius-6 patch of grid cells around the barycentric point $(i_0, j_0) = (\mathrm{round}(n/3), \mathrm{round}(n/3))$ within one icosahedral face therefore suffices to build every triangle needed to evaluate definition 25.10's plane-projected defect at the centroid triangle and its immediate neighbors, at $O(1)$ cost in $n$.

**Proposition 25.57** (Validated exactly against the full mesh)**.** *`scripts/m_centroid_localized_search.py` searches this patch for its own local maximum defect, at $n = 51$ and $n = 99$ — the two odd multiples of 3 nearest section 25.7.37's own tested frequencies for which section 25.7.37's full-mesh routine (`m_n_argmax_locus.py`) can independently be re-run in this pass. The full mesh gives $M(51) = 1.20565434$ with argmax defect $0.01668726°$ at ring 16 (one ring short of $n/3 = 17$), and $M(99) = 1.20632462$ with argmax defect $0.00443094°$ at ring 32 (one ring short of $n/3 = 33$); the local patch, searching a radius-6 neighborhood of the naive centroid triangle rather than assuming the centroid triangle itself is the maximizer, independently finds $0.016687257\ldots°$ and $0.004430944\ldots°$ respectively — agreement with the full-mesh value to every digit the full-mesh routine itself reports, at the correct off-centroid ring in both cases. This is the same validation discipline section 25.7.29's own construction applied to the corner locus, carried out here for the first time at the centroid locus.*

**Key Result**

*[Pushed to $n = 10{,}000$ — twenty-five times section 25.7.38's own full-mesh ceiling — $M_\infty$ is pinned down an order of magnitude tighter, and agrees with the closed form within the fit's own noise] Repeating the same radius-6 search at $n = 201, 300, 402, 1000, 3000, 10000$ (each under a tenth of a second; no full mesh is ever built) gives $M(n) = 1.2065082087$, $1.2065406225, 1.2065523216, 1.2065657063, 1.2065667637, 1.2065670252$ respectively — monotonically increasing at every step, consistent with proposition 25.55's own $n = 50$-to-400 values $(1.206039, \ldots, 1.206559)$ over the range where the two overlap. Fitting $M(n) = M_\infty - c/n^2$ (proposition 25.55's own model) by ordinary least squares over these six points gives*

$$M_\infty^{(\text{local fit})} = 1.2065673102, \qquad c = 2.3907,$$

*with residuals no larger than $7.9 \times 10^{-7}$ (at $n = 1000$) — an order of magnitude tighter than proposition 25.55's own $8.2 \times 10^{-6}$ residual from five points reaching only $n = 400$. Comparing against section 25.7.38's closed form,*

$$M_\infty^{(\text{area})} = \frac{15\left(7\sqrt{3} - 3\sqrt{15}\right)}{2\pi} = 1.2065670269\ldots,$$

*the two differ by $2.8 \times 10^{-7}$ — smaller than this fit's own maximum residual, the identical standard proposition 25.55 itself used to call its $n = 400$-ceiling agreement "indistinguishable from it given the data." The present fit reaches that same conclusion from data twenty-five times further out in $n$, with an order-of-magnitude tighter residual, and is an independent construction (a local patch, not a rebuilt full mesh) rather than a rerun of the same routine at larger scale.*

**Epistemic Notice**

Pushing further — $n$ = 30,000, 100,000, 300,000, 1,000,000 — does *not* continue this improvement: the search returns $M$ = 1.2065675906, 1.2065752240, 1.2067024461, 1.2086107790, a departure that grows and eventually accelerates, the same qualitative shape section 25.7.29's own extreme-$n$ breakdown showed for $\kappa(n)$, but traced here to a different mechanism. A direct diagnostic (evaluating $2\pi$ minus the three plane-projected angles at the single, fixed centroid triangle $(i_0, j_0) = (\text{round}(n/3), \text{round}(n/3))$, without searching its neighbors) shows that triangle's own defect remaining numerically sane — $4.34 \times 10^{-9}$, $4.83 \times 10^{-10}$, $4.35 \times 10^{-11}$ degrees at $n$ = 100,000, 300,000, 1,000,000 respectively, tracking the expected $O(1/n^2)$ decay reasonably even at $n = 10^6$ — because the underlying angle sum agrees with $2\pi$ to 6.283185307…, i.e. to the 10th significant figure or beyond at every $n$ tested, leaving a residual at the edge of float64's own absolute cancellation floor (of order $1 \times 10^{-14}$ to $1 \times 10^{-13}$ degrees in this construction) rather than losing the exponent's shape entirely. The *search* degrades far earlier than that single fixed evaluation does, for a selection reason rather than a per-evaluation one: proposition 25.57's own full-mesh cross-check already found 60 and more near-tied candidate faces at the true maximum (section 25.7.37's own 80-tied-at-$n$=100, 2060-tied-at-$n$=200 observation, for the same reason), so once the *true* gap between the top handful of candidates falls below that same $1 \times 10^{-14}$-degree floor — which section 25.7.40's own decaying defect values show happening somewhere between $n$ = 10,000 and $n$ = 30,000 — the reported "maximum" is decided by which candidate's cancellation noise happens to land positive, a selection effect that biases the search upward rather than merely adding symmetric noise to a single value. This is a genuinely different failure mode from section 25.7.29's own ceiling (there, the barycentric weight $(n-2)/n$ itself collapsing to 1 in float64 near $n \sim 2 \times 10^5$–$3 \times 10^5$, a problem with the *points*; here, a search over many near-degenerate candidates amplifying an already-present cancellation floor in the *angle sum*, a problem with the *selection*, present already by $n \sim 3 \times 10^4$) — but both are, in the same spirit, floating-point ceilings intrinsic to a specific $O(1)$-per-$n$ construction rather than failures of the underlying mathematics, and both would in principle yield to `mpmath`-style arbitrary-precision arithmetic not attempted in this pass.

**Todo — verify primary source**

Section 25.7.40's $2.8 \times 10^{-7}$ agreement, while a genuine tightening, is not the proof section 25.7.38's own closing todobox asked for, and this section does not attempt the local expansion of $\delta(v) - K \cdot A(v)$ near $\alpha = \beta = \gamma = 1/3$ that todobox specified as what a genuine proof would require; it substitutes a numerically sharper anchor ($n$ = 10,000 in place of $n$ = 400) for the same style of fit-based evidence already in hand, not a different kind of evidence. Two further, narrower items are newly opened by this section's own diagnostic rather than closed by it: first, whether the single-fixed-triangle evaluation's own float64 cancellation floor (of order $1 \times 10^{-14}$ degrees, identified but not further reduced here) could itself be pushed down with `mpmath` to extend either the search or the fixed-triangle fit past $n \sim 10^4$–$10^5$ without the search's own selection-bias failure mode, is untried; second, whether the same selection-bias mechanism — rather than the weight-collapse this chapter originally diagnosed for $\kappa_\infty^{\text{SSB}}$ — also contributes to section 25.7.29's own $n \gtrsim 3 \times 10^5$ ceiling, given that the pentagon-corner locus is likewise exactly fivefold-degenerate rather than a unique maximum, has not been checked. Unchanged, as always: whether $M(n)$ and $\text{CV}(n)$ converge at all is still asserted only on the strength of a fit, not proved; $\text{CV}(n)$ itself still has no localized construction of any kind, for the reason section 25.7.37's own todobox already gave; and $\kappa_\infty^{\text{SSB}}$'s closed form remains exactly as open as ever.

### 25.7.41 The fixed-triangle evaluation, in arbitrary precision: the float64 floor removed, not merely lowered

The preceding todobox's first item is untried no longer. `scripts/m_centroid_mpmath.py` reimplements section 25.7.40's *fixed*-centroid-triangle evaluation — the single defect computation at $(i_0, j_0) = (n/3, n/3)$, not the radius-6 search around it — entirely in `mpmath` arbitrary-precision arithmetic: the icosahedron's vertices, the barycentric interpolation onto

the local patch, the plane-fit normal (via `mpmath`'s symmetric eigensolver `eigsy` on the $3 \times 3$ covariance matrix of each hexagonal star, in place of `numpy`'s SVD) and the resulting angle are all carried at working precision throughout, with no float64 arithmetic anywhere in the path. Deliberately excluded from this treatment is the radius-6 *search*: section 25.7.40's own disclaimerbox diagnosed that search's breakdown as a selection effect among dozens of near-tied candidate faces rather than a per-evaluation precision floor, so extending it correctly would require resolving those near-ties, not merely computing each candidate more precisely — a distinct undertaking from the one attempted here, and the second item the preceding todobox itself left open, still open below.

**Proposition 25.58** (Validated against every float64 value on record)**.** *At working precision* 30 *decimal digits,* `m_centroid_mpmath.py` *reproduces proposition 25.57's float64 fixed-triangle values at* $n = 201, 300, 402$ *and section 25.7.40's value at the nearest tested multiple of* 3 *to every digit either float64 computation reports (*$M(201) = 1.20650820861\ldots$*,* $M(300) = 1.20654062234\ldots$*,* $M(402) = 1.20655232154\ldots$*, matching section 25.7.40's* 1.2065082087*,* 1.2065406225*,* 1.2065523216 *within float64's own last-digit rounding).*

**Key Result**

*[Twelve orders of magnitude past the float64 ceiling: the residual is exactly the leading $1/n^2$ term, coefficient stable to ten figures] Raising working precision — on an adaptive schedule rather than a single fixed digit count, in the same spirit as section 25.7.29's own precision schedule (proposition 25.47): 50 decimal digits for $n = 999$ through 99,999, increased in small steps to 64 digits by $n \approx 10^{12}$ — and evaluating at $n = 999, 9999, 99999, 999999, 9999999, 99999999, 999999999, 9999999999$ (the largest, $n \approx 10^{12}$, a factor of $2.5 \times 10^8$ past section 25.7.40's own $n = 10{,}000$ float64-search ceiling, and one that the float64 fixed-triangle evaluation's own $1 \times 10^{-14}$-degree cancellation floor made entirely inaccessible), each evaluation still completes in well under a tenth of a second — no slower than the float64 version, since the construction touches only $O(1)$ points regardless of $n$. The adaptive schedule is not a stylistic choice but a checked necessity: a single fixed working precision of 55 digits, adequate through $n = 999{,}999{,}999$, was tried first at $n \approx 10^{12}$ as well and found to fall short there — the reported coefficient drifts to $-2.376482874$, visibly off in the eighth significant figure from the value every smaller $n$ agrees on; raising to 64 digits at that one frequency removes the drift and restores $-2.376482717$. This is itself informative: it confirms the earlier apparent stability at fixed precision was not yet being tested to its own breaking point, and that the coefficient's true stability has to be established digit-budget by digit-budget as $n$ grows, exactly as section 25.7.29's own 60-to-80-digit schedule already had to for $\kappa(n)$, rather than assumed from a single precision choice applied uniformly. Writing $\Delta(n) := M(n) - M_\infty^{(\mathrm{area})}$ against section 25.7.38's closed form $M_\infty^{(\mathrm{area})} = 1.206\,567\,026\,876\,591\,348\,921\,031\,915\,97$ (computed here to 30 digits from the same expression), the values obtained — each at working precision sufficient that a further increase leaves every displayed digit unchanged — are:*

| $n$ | $\Delta(n)$ | $\Delta(n) \cdot n^2$ |
|---|---|---|
| 201 | $-5.881\,826\,715 \times 10^{-5}$ | $-2.376316811$ |
| 999 | $-2.381\,236\,091 \times 10^{-6}$ | $-2.376476$ |
| 9999 | $-2.376\,958\,018 \times 10^{-8}$ | $-2.37648265$ |
| 99999 | $-2.376\,530\,247 \times 10^{-10}$ | $-2.376482716$ |
| 999999 | $-2.376\,487\,47 \times 10^{-12}$ | $-2.376482717$ |
| 9999999 | $-2.376\,483\,192 \times 10^{-14}$ | $-2.376482717$ |
| 99999999 | $-2.376\,482\,764 \times 10^{-16}$ | $-2.376482717$ |
| 999999999 | $-2.376\,482\,722 \times 10^{-18}$ | $-2.376482717$ |
| 999999999999 | $-2.376\,482\,717 \times 10^{-24}$ | $-2.376482717$ |

*$\Delta(n)$ is negative and monotonically shrinking at every step, and the rescaled quantity $\Delta(n) \cdot n^2$ — not merely bounded, as section 25.7.40's five-point fit over $n = 201$ to 10,000 already suggested, but genuinely convergent — settles to $-2.376482717$ and does not move in the ninth significant figure across the last five entries, spanning $n = 99{,}999$ to $n \approx 10^{12}$, seven further orders of magnitude in $n$ with the coefficient unchanged.*

**Epistemic Notice**

What section 25.7.41 establishes and does not establish should be stated with the same care proposition 25.54's own disclaimerbox used. This is not the analytic proof section 25.7.38's original todobox asked for: no local expansion of $\delta(v) - K \cdot A(v)$ near $\alpha = \beta = \gamma = 1/3$ is derived here, and the coefficient $-2.376482717\ldots$ is read off numerically rather than computed in closed form from such an expansion (nor is it yet identified with any of this chapter's other closed-form constants, e.g. by inspection against $\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi$ combinations of the kind section 25.7.38 itself found). What it does establish is considerably more than section 25.7.40's own $2.8 \times 10^{-7}$ figure claimed: removing the float64 floor does not merely shrink the gap between $M(n)$ and the closed form further, as more precision applied to a genuinely open question might or might not do — it shows the gap tracking a single, stable $c/n^2$ law across twelve orders of magnitude in $n$ with no sign of the coefficient itself drifting, which is exactly the signature a true asymptotic expansion produces and exactly what section 25.7.25's earlier, ultimately unsuccessful test of the analogous hypothesis for $\kappa_\infty^{\mathrm{SSB}}$ did *not* show (there, section 25.7.24's 4% gap did not close under refinement at all). Read alongside that contrast, section 25.7.41 is the strongest evidence this volume has assembled anywhere for one of its area-based closed-form conjectures being exact rather than a leading-order approximation — short of a proof, but no longer just a numerically plausible coincidence at the float64 precision this chapter's earlier passes were confined to.

**Todo — verify primary source**

Section 25.7.41 closes the first of the two items the preceding todobox left open — `mpmath` does remove, not merely lower, the fixed-triangle evaluation's precision ceiling — but leaves three genuinely open. First, and most substantively: an analytic derivation of the coefficient $c \approx 2.376482717$, either from the local expansion of $\delta(v) - K \cdot A(v)$ that section 25.7.38's original todobox specified, or by recognizing $2.376482717\ldots$ in closed form (a short numerical search against low-degree combinations of $\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi$ was not attempted in this pass), would upgrade section 25.7.41 from very strong numerical evidence to an actual proof, and remains exactly as absent as before this section. Second, the search-based construction's own selection-bias ceiling (section 25.7.40's disclaimerbox) is untouched by this section, which deliberately worked around it rather than resolving it; doing so in `mpmath` would need the near-tied candidates themselves resolved at high precision, not merely each one evaluated more precisely, since the float64 selection failure is about *which* candidate is largest, not how accurately any one candidate's own defect is known. Third, the preceding todobox's own second item — whether an analogous selection-bias mechanism, rather than the weight-collapse originally diagnosed, contributes to section 25.7.29's $\kappa_\infty^{\mathrm{SSB}}$ ceiling — is unchanged by anything in this section, which touched only the $M(n)$ construction; `mpmath` applied to section 25.7.29's own fixed-point (rather than search) evaluation, in the same spirit as this section, is a natural next step but was not attempted here. Unchanged, finally: CV($n$)'s convergence still rests only on sections 25.7.6 and 25.7.32's fits, with no localized construction of any kind proposed for it; and $\kappa_\infty^{\mathrm{SSB}}$'s own closed form remains, on the evidence section 25.7.24 already found and section 25.7.41's own contrasting cleanliness now underscores by comparison, genuinely short by a real 4% rather than merely awaiting more precision.

### 25.7.42 The coefficient in closed form: $c = 405\sqrt{3}(47 - 21\sqrt{5})/(4\pi)$, to 69 stable digits

The preceding todobox's first and most substantive item — recognizing $c \approx 2.376482717\ldots$ in closed form, or deriving it analytically from the local expansion of $\delta(v) - K \cdot A(v)$ — is closed here, by the numerically-driven route the todobox itself flagged as untried rather than by the analytic route it also named. `scripts/m_centroid_coefficient_closed_form.py` exploits the same fact section 25.7.41 already relied on — `defect_at_centroid_triangle`'s cost is $O(1)$ in $n$, so pushing $n$ far larger costs nothing beyond the working precision needed to survive cancellation — but pushes it much further: to $n \approx 10^{30}$ and $n \approx 10^{80}$, at 200 and 550 decimal digits of working precision respectively, rather than to $n \approx 10^{12}$ at 64 digits as section 25.7.41 did. At these frequencies $\Delta(n)$'s own $O(1/n^4)$ and higher sub-leading

terms are suppressed by well over a hundred further orders of magnitude past the leading $c/n^2$ term, so $n^2\big(M_\infty^{(\mathrm{area})} - M(n)\big)$ recovers $c$ itself directly, to as many stable digits as the working precision allows, with no fitting or extrapolation across multiple $n$ needed — unlike every earlier numerical estimate of $c$ in this chapter, including section 25.7.41's own.

**Proposition 25.59** (Two independent runs agree to sixty digits)**.** *The $n \approx 10^{30}$ (prec 200) and $n \approx 10^{80}$ (prec 550) runs give*

$$c = 2.37648271692499244699390015120506913236495640335974619801 92\ldots$$

*agreeing digit-for-digit for sixty places before the shorter run's own working precision is exhausted (difference $-6.703\,33 \times 10^{-60}$, consistent with the 200-digit run's own precision floor, not with any discrepancy in the underlying quantity).*

**Key Result**

*[The coefficient is $405\sqrt{3}(47 - 21\sqrt{5})/(4\pi)$, matched to 69 digits] An* `mpmath.pslq` *search over integer relations between $c$ and a 20-term basis of low-degree combinations of $\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi$ (the same four constants section 25.7.38's own $M_\infty^{(\mathrm{area})}$ is built from, together with $\pi^{\pm1}, \pi^{\pm2}$ multiples of each) returns, from nineteen equivalent presentations of the same relation, the integer identity $-4c + 19035 \cdot (\sqrt{3}/\pi) - 8505 \cdot (\sqrt{15}/\pi) = 0$. Since $19035 = 3^4 \cdot 5 \cdot 47$ and $8505 = 3^5 \cdot 5 \cdot 7$ share the common factor $405 = 3^4 \cdot 5$, and $\sqrt{15} = \sqrt{3}\sqrt{5}$, this reduces to*

$$c \;=\; \frac{405\sqrt{3}\,(47 - 21\sqrt{5})}{4\pi}.$$

*Evaluated at 70 decimal digits of working precision, this closed form agrees with* both *of proposition 25.59's independently-computed values of $c$ to every digit each is good for: difference $6.703\,33 \times 10^{-60}$ against the $n \approx 10^{30}$ run and $-1.984\,52 \times 10^{-69}$ against the $n \approx 10^{80}$ run — i.e. to the full sixty-nine digits the longer run resolves, with no sign of divergence at the digit where either run's own working precision runs out.*

**Epistemic Notice**

What section 25.7.42 establishes and does not establish should be read with the same care every other closed-form claim in this section has been given. This is numerical identification, not algebraic derivation: no expansion of $\delta(v) - K \cdot A(v)$ near $\alpha = \beta = \gamma = 1/3$ to next order in $1/n$ is carried out here, so the mechanism producing the specific integers 47 and 21 (neither of which is a small icosahedral-symmetry number in the way 7 and 3 in $M_\infty^{(\mathrm{area})}$'s own numerator plausibly are) remains unexplained by anything in this section. What distinguishes this identification from a merely suggestive numerical coincidence is the same thing that distinguished section 25.7.38's own eventual acceptance: an agreement of *tens of digits*, obtained from *two independently-computed* high-precision values of $c$ at two different, widely-separated $n$, against a closed form built from a small integer relation ($|-4|, 19035, 8505$, none exceeding five digits) found by an unbiased search over a modest fixed basis rather than reverse-engineered to fit. A spurious coincidence of this precision, from integer coefficients this small against a basis this size, is not a realistic alternative explanation; PSLQ relations of this quality are the standard form of evidence this kind of constant-recognition problem produces, short of an actual proof. It is that standard of evidence, and not more, that is claimed here.

**Todo — verify primary source**

Section 25.7.42 closes the preceding todobox's first and most substantive item; two remain, both unchanged by anything in this section, which touched only the numerical value of $c$ and not either open construction question. First, the search-based construction's own selection-bias ceiling (section 25.7.40's disclaimerbox) is still untouched: resolving it would need `mpmath` applied to the near-tied *candidates* themselves at high precision, a distinct undertaking from the single fixed-triangle evaluation this section and section 25.7.41 both worked with instead. Second, whether the same selection-bias mechanism (rather than the weight-collapse originally diagnosed) contributes to section 25.7.29's own $\kappa_\infty^{\mathrm{SSB}}$ ceiling remains unchecked, as does the narrower question of whether $\kappa_\infty^{\mathrm{SSB}}$ itself — on the evidence section 25.7.24 already found, a genuine 4% short of its own area-based closed-form candidate rather than a precision artifact — might, like $M_\infty$ here, eventually admit a corrected or entirely different closed form of its own; section 25.7.42's technique (push $n$ far past any fitting range, then PSLQ the single resulting high-precision value directly) is directly transferable to that question but was not attempted here, since $\kappa_\infty^{\mathrm{SSB}}$'s own localized fixed-point evaluation (as opposed to section 25.7.29's corner-search construction) does not yet exist in `mpmath` form, unlike $M(n)$'s. Unchanged, finally, exactly as before: CV($n$)'s convergence rests only on a fit, with no localized construction proposed for it at all.

### 25.7.43 The search itself, in arbitrary precision: the selection-bias hypothesis tested directly

The preceding todobox's first item — whether the radius-6 search's own float64 breakdown, past roughly $n \sim 3 \times 10^5$, is a genuine feature of $M(n)$ or an artifact of `m_centroid_localized_search.py`'s float64 candidate comparison — is taken up directly here, rather than worked around as sections 25.7.41 and 25.7.42 both did by fixing the evaluation point in advance. `scripts/m_centroid_search_mpmath.py` reimplements the search itself — not merely the single centroid-triangle evaluation — entirely in `mpmath`: the icosahedron, the radius-6 local patch construction, and every candidate face's plane-projected defect are all carried at working precision throughout, and the candidates are then ranked by their `mpmath` values rather than by float64 comparison.

Validation first, on the same discipline as every other script in this appendix: at $n = 201, 300, 402$, where float64 is not yet in its failure regime, the `mpmath` search (at 30 working digits) reproduces section 25.7.40's own published float64 values — $M(201) = 1.2065082086\ldots$, $M(300) = 1.2065406225\ldots$, $M(402) = 1.2065523215\ldots$ — to every digit float64 itself reports.

**Key Result**

*[The float64 divergence at large $n$ is a pure selection artifact, not a signal about $M(n)$] Pushed through and past section 25.7.40's own breakdown range, at an adaptive precision schedule (40 digits through $n = 10^5$, 50 digits at $n = 3 \times 10^5$ and $n = 10^6$), the* `mpmath` *search converges cleanly:*

| $n$ | $M(n)$ (`mpmath search`) |
|---|---|
| 30 000 | 1.206567024236… |
| 100 000 | 1.206567026744… |
| 300 000 | 1.206567026850… |
| 1 000 000 | 1.206567026875… |

*against section 25.7.38's closed form $M_\infty^{(\text{area})} = 1.20656702687659\ldots$ — each value strictly closer than the last, with no sign of the drift that overtook the float64 search at exactly these same frequencies. The float64 script, over the identical range, gives $M(30{,}000) = 1.2065675906\ldots$ (already six parts in $10^6$ off the true value at this $n$) worsening to $M(1{,}000{,}000) = 1.2086107790\ldots$ (a departure of nearly $2 \times 10^{-3}$, two orders of magnitude larger than the leading $c/n^2$ term at that frequency). Section 25.7.40's disclaimerbox posed this as an open question between two hypotheses — a genuine feature of $M(n)$'s search-based construction, or a float64 selection artifact — and this is now settled in favor of the second: the search's own arithmetic, not the quantity it is searching for, was unstable.*

The mechanism is confirmed directly, not merely inferred from the clean convergence above. At $n = 300{,}000$ — the frequency at which section 25.7.40's float64 search first visibly departs from the $c/n^2$ trend — the `mpmath` search's own ranked candidate list, computed at 90 working digits, shows two of the radius-6 patch's 242 candidate faces agreeing to every one of the 80 decimal digits both are printed to; a third and fourth pair agree to the same depth one rank lower. This is not a numerical near-miss that finer precision would eventually resolve: at 90 digits, the top two candidates' defects differ by less than $1 \times 10^{-80}$ degrees, i.e. by exactly zero at the precision tested, which is the signature of an exact degeneracy enforced by the icosahedron's own rotational symmetry (distinct radius-6 candidate faces sitting at symmetry-equivalent positions relative to the centroid locus) rather than of two merely close but distinct real numbers. Float64's $\sim$ 15-digit working precision cannot resolve a tie of this kind at all once the surrounding defect values themselves have shrunk, at large $n$, to within float64's own representable range of that degeneracy — it does not merely round the winning margin poorly, it has no margin left to round. This is exactly section 25.7.40's disclaimerbox's own selection-bias hypothesis, now directly observed rather than only inferred from the shape of the divergence.

### 25.7.44 Does the same mechanism reach $\kappa_\infty^{\text{SSB}}$? A structural answer, not a new computation

The preceding todobox's second item asked whether an analogous selection-bias mechanism, rather than the barycentric-weight collapse originally diagnosed, also contributes to section 25.7.29's own $n \gtrsim 3 \times 10^5$ ceiling for $\kappa_\infty^{\text{SSB}}$. This does not require a new large-$n$ push of the kind section 25.7.43 just carried out for $M(n)$: inspection of `scripts/kappa_localized_mpmath.py`'s own `kappa_of_n` shows the two constructions are not analogous in the one respect that matters.

`kappa_of_n` computes $\kappa(n)$ by calling `defect_at_corner` once for each of the icosahedron's twelve vertices, and within each call, `local_patch` identifies exactly five pentagon-corner triangles — `corner_tri_idx`, built directly from the five faces incident to that vertex in the base icosahedron, not searched for — and evaluates the defect at each of those five known positions. The resulting sixty values are simply collected (`all_d.extend(…)`); nothing in the function compares candidates against one another or picks a maximum among

near-tied contenders. Section 25.7.43's degeneracy, by contrast, arose specifically from `m_centroid_localized_search.py`'s `argmax_defect_near_centroid`: a radius-6 patch of *combinatorially unknown-position* candidate faces, ranked by computed defect to find whichever one is largest. Selection bias, in the sense section 25.7.40's disclaimerbox used the term, is specifically the failure of that ranking step under float64 precision when several candidates are nearly (or exactly) tied; a construction that never ranks candidates has no step for that failure to occur in.

**Key Result**

*[Structurally closed: $\kappa_\infty^{\mathrm{SSB}}$'s known-position construction has no selection step] $\kappa_\infty^{\mathrm{SSB}}$'s pentagon-corner defects are read off at sixty combinatorially exact, individually enumerated positions — twelve icosahedral vertices × five corner triangles each, fixed by the base icosahedron's own connectivity and independent of $n$ — never chosen by comparing candidates of a priori unknown position. Section 25.7.29's own $n \gtrsim 3 \times 10^5$ ceiling therefore cannot be a selection-bias effect of the kind section 25.7.43 just found and confirmed for $M(n)$'s search; the original diagnosis — collapse of the barycentric weight $(n-2)/n$ to 1.0 in float64, named when that ceiling was first reported — stands as it was, not as a partial account awaiting this correction.*

**Todo — verify primary source**

Sections 25.7.43 and 25.7.44 close both items the preceding todobox left open, but on asymmetric evidence: the first by a genuine new computation (the radius-6 search rerun end to end in `mpmath`, validated against the published float64 table before being trusted past it), the second by inspection of code already present in this volume, not by any new numerical push. Nothing here constitutes progress on $\kappa_\infty^{\mathrm{SSB}}$'s own closed form, which section 25.7.24 already found to fall a genuine 4% short of its area-based candidate — a real gap, on that evidence, rather than a precision ceiling of any kind, selection-bias or otherwise. What section 25.7.44 does newly establish is that section 25.7.42's own push-then-PSLQ technique remains available, unblocked by this pass's own question, as a route to test whether $\kappa_\infty^{\mathrm{SSB}}$ admits a corrected or altogether different closed form of its own — conditional, as section 25.7.42's own todobox already noted, on an `mpmath` *fixed-point* (not search) evaluation of $\kappa(n)$ existing at arbitrarily large $n$, which `kappa_localized_mpmath.py` in its present form does not yet provide (it validates against $n = 8, 10, 20$ only) and which remains a genuinely open, concretely scoped next step rather than a claim made here.

### 25.7.45 The fixed-point push, taken up: a real bug found on the way, and $\kappa_\infty^{\mathrm{SSB}}$ pushed to roughly fifty stable digits

The preceding todobox's own concretely-scoped next step — an `mpmath` fixed-point evaluation of $\kappa(n)$ at arbitrarily large $n$, which `kappa_localized_mpmath.py` did not yet provide — is taken up directly here by `scripts/kappa_mpmath_push_and_pslq.py`, on the same validate-then-push discipline as every other script in this appendix.

Validation first reproduces `kappa_localized.py`'s own published float64 values through $n = 100{,}000$ exactly. Pushing past that point surfaced a genuine defect in `kappa_localized_mpmath.py` itself, not in the mathematics: `local_patch`'s point-deduplication step keyed each grid point by a string rounded to a *fixed* 25 significant digits, regardless of the working precision (`mp.mp.dps`) actually in force. Since two adjacent grid points at frequency $n$ can be separated by as little as $O(1/n)$ in relative terms, this key silently merges genuinely distinct points once $n$ exceeds roughly $10^{25}$ — observed directly, before the fix, as an `order_star` lookup failure at $n = 10^{30}$ (a corrupted local patch topology, not a precision artifact of the angle computation itself). The fix ties the key's own precision to the working precision (`mp.mp.dps + 5` digits) rather than a fixed constant, the same discipline the script's adaptive precision schedule already applies elsewhere. This

is recorded as a bug fix to existing volume code, in the same spirit section 25.7.29's own two float64 precision issues (the arccos instability, fixed; the barycentric-weight collapse, diagnosed and left unfixed) were recorded when first found.

**Key Result**

*[$\kappa_\infty^{\mathrm{SSB}}$ to roughly fifty stable digits] With the key fixed, $\kappa(n)$ is stable in $n$ at $n = 10^{20}$, $10^{30}$, $10^{80}$, and $10^{150}$ (working precision 80 through 350 digits, each confirmed independently stable under a further precision increase at fixed $n$, ruling out a working-precision floor rather than a true $n$-dependence): the $n = 10^{80}$ and $n = 10^{150}$ values agree to the full fifty-one digits both are printed to,*

$$\kappa_\infty^{\mathrm{SSB}} \approx 0.60282350203376711628173127137708096916153250194790\ldots,$$

*extending section 25.7.28's own five-significant-figure estimate (0.60282) by more than forty further stable digits. The residual's own decay rate is not fitted here — this section pushes a single $n$ far past any plausible fitting range, exactly as section 25.7.42's own technique did for $M_\infty$'s coefficient $c$, rather than extrapolating across several — but comparing the $n = 10^{30}$ and $n = 10^{80}$ values (agreeing to only about thirty digits, against fifty between $10^{80}$ and $10^{150}$) is consistent with a leading correction term that shrinks distinctly slower than $M_\infty$'s own clean $c/n^2$ law; no exponent is fitted or claimed from this observation alone.*

**Key Result**

*[No closed form found, against the exact basis that found $M_\infty$'s] Section 25.7.45's value, screened by `mpmath.pslq` at 110 working digits against precisely the same seventeen-term basis ($1$, $\sqrt{3}, \sqrt{5}, \sqrt{15}$, and each divided by, multiplied by, or divided by the square of $\pi$) that returned section 25.7.42's integer relation for $M_\infty$'s coefficient $c$, returns no relation at all. This is a materially stronger negative result than section 25.7.28's own closed-form screen (eighteen icosahedrally-motivated candidates and every low-denominator rational, tested at ordinary numerical precision): the same basis, the same search technique, and over a hundred correct digits to test it against, and still nothing — ruling out, specifically, the possibility that $\kappa_\infty^{\mathrm{SSB}}$ shares $M_\infty^{(\mathrm{area})}$'s own algebraic ingredients ($\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi$) in any low-degree combination, not merely that a handful of hand-picked candidates failed to match at low precision.*

**Todo — verify primary source**

Section 25.7.45 answer the preceding todobox's own request, but the answer is a sharpened negative, not a closure: $\kappa_\infty^{\mathrm{SSB}}$ still has no known closed form, and this section rules out one specific, previously-untested family of candidates (the $M_\infty$-style basis) rather than finding one. Three items follow directly from this pass's own work and are newly opened, or newly sharpened, by it. First, the observed slow-down in convergence rate relative to $M_\infty$ (section 25.7.45's closing remark) is only a qualitative impression from three data points, not a fit; a genuine Richardson-style extrapolation across several large, widely-spaced $n$ (in the manner `highprec_kappa.py` already applies to the GEO-role constants, section 25.7.13's own technique) would both pin down the exponent and, if it is smaller than 2, explain why $\kappa_\infty^{\mathrm{SSB}}$ has resisted exactly this kind of high-precision attack for longer than $M_\infty$ did — this is not attempted here. Second, section 25.7.45's basis, while identical to the one that worked for $M_\infty$, is one basis among many; a wider or differently-motivated basis (built, for instance, from the icosahedral central angle $\arccos(1/\sqrt{5})$ already screened, unsuccessfully, in `highprec_kappa.py`, or from the golden ratio $\varphi$ directly rather than only through $\sqrt{5}$) has not been tried at this pass's precision and remains open. Third, and most substantively: section 25.7.24's own 4% area-hypothesis residual is entirely untouched by anything in this section, which sharpens only the closed-form search, not the mechanistic question of what that residual represents — the volume's standing position, that the residual is tied to fivefold combinatorial valence under chord-based discretization and addressable only by the still-unobtained Feinberg–Friedberg–Lee–Ren machinery (section 25.7.21), stands exactly as it did before this section.

### 25.7.46 A second basis, from the golden ratio and the icosahedral central angle directly

The preceding todobox's second item — whether a basis built directly from $\varphi = (1+\sqrt{5})/2$ and from the icosahedral central edge-angle $\theta_e = \arccos(1/\sqrt{5})$, rather than only from $\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi$, might recognize $\kappa_\infty^{\mathrm{SSB}}$ — is taken up here by `scripts/kappa_phi_theta_basis_pslq.py`. $\theta_e$ is the angle subtended at the icosahedron's centre by two adjacent vertices; it is the natural candidate for a second, differently-motivated basis because, unlike $\sqrt{5}$ alone, it enters the construction through the actual solid angle geometry of the polyhedron rather than only through the golden ratio's algebra.

Building this basis correctly took one false start worth recording rather than silently discarding. A first attempt included, at a single scale, three or more of $\{\varphi, 1/\varphi, \varphi^2, \sqrt{5}\}$ together; since $\varphi^2 = \varphi + 1$, $1/\varphi = \varphi - 1$, and $\sqrt{5} = 2\varphi - 1$ are each a *rational* linear combination of $\{1, \varphi\}$, any such set is linearly dependent over $\mathbb{Q}$ independently of $\kappa_\infty^{\mathrm{SSB}}$ entirely. `mpmath.pslq` duly returned a relation each time — but with coefficient exactly zero on $\kappa_\infty^{\mathrm{SSB}}$ itself: the two discarded outputs were the trivial identities $\theta_e(1 - \varphi + 1/\varphi) = 0$ and $2\varphi - \sqrt{5} - 1 = 0$, both true of the basis alone and telling nothing about $\kappa_\infty^{\mathrm{SSB}}$. This is the same failure mode a screen against a redundant basis always risks, and the reason every closed-form screen elsewhere in this appendix has been built from a basis checked, at least informally, for this kind of in-family redundancy; this section's own first pass simply failed to check it before running `pslq`, and the check is recorded here so a future pass building a further basis does not repeat it silently. The corrected basis keeps at most $\{1, \varphi\}$ from the golden-ratio field at each scale ($1$, $1/\pi$, $1/\pi^2$) and $\theta_e$ at only two scales ($1$, $1/\pi$), alongside $\pi$, $\sqrt{2}$, and $\sqrt{3}$ at scale 1 and $1/\pi$ — thirteen basis terms in total, none a rational combination of any other.

**Key Result**

*[No relation found against the golden-ratio/central-angle basis either]* $\kappa_\infty^{\mathrm{SSB}}$ *is recomputed directly from* `kappa_mpmath_push_and_pslq.py`*'s own* `kappa_at` *function at* $n = 10^{80}$*, working precision* 220 *digits — not copied from section 25.7.45's own printed fifty-one digits, so this section's screen rests on an independently-regenerated value, and the recomputation agrees with that printed value over its own full length, a small consistency check in its own right. Screened by* `mpmath.pslq` *against the corrected thirteen-term* $\{1, \varphi, \theta_e, \pi, \sqrt{2}, \sqrt{3}\}$*-basis above, at* 220 *working digits, the search returns no relation at all.*

This is a second, differently-grounded negative result, not a repetition of section 25.7.45 under a new name: that section ruled out $M_\infty^{(\mathrm{area})}$'s own algebraic ingredients; this one rules out, specifically, any low-degree combination of the golden ratio and the icosahedron's own central edge-angle, tested at the same order of precision. Between the two, the two most geometrically motivated families of candidate closed forms available to this volume — one from the flat-chord area Jacobian that gave $M_\infty$ its exact value, one from the icosahedron's own defining angles — have now both been screened at high precision and both come back empty.

**Todo — verify primary source**

Section 25.7.46 closes the preceding todobox's second item, on the same terms section 25.7.45 closed its first: a sharpened negative, not a closure of $\kappa_\infty^{\mathrm{SSB}}$'s own closed-form question, which remains genuinely open. Two paths forward are more clearly delimited by having eliminated both of these families. First, a basis built from neither family — for instance, from $\arccos(1/3)$ (the icosahedron's own dihedral angle, distinct from $\theta_e$) or from values of the digamma or elliptic-integral functions at rational arguments, the kind of constant that appears in exact Regge-calculus defect sums for related polyhedra — has not been tried and is not obviously motivated by anything already in this chapter's own construction, unlike $\theta_e$, which at least enters the geometry directly. Second, and arguably more informative than a further basis search: section 25.7.45's own closing remark, that $\kappa_\infty^{\mathrm{SSB}}$'s residual appears to shrink slower than $M_\infty$'s clean $c/n^2$ law, is still unfitted; if a genuine Richardson extrapolation across the already-computed $n = 10^{20}, 10^{30}, 10^{80}, 10^{150}$ points found a non-integer or non-simple exponent, that alone would be evidence against a closed form of the same algebraic character as $M_\infty^{(\mathrm{area})}$'s existing at all, independent of which basis is tried against it — a structural argument for why the closed-form search keeps failing, rather than one more attempt at the search itself. Neither is undertaken here.

### 25.7.47 A third basis, from the flat-chord construction's own normalization step

The preceding todobox's first item is taken up here, but not along the specific line it suggested ($\arccos(1/3)$ or elliptic-integral values): `scripts/kappa_chord_norm_basis_pslq.py` instead builds a third basis from the one algebraic quantity that enters SSB's own construction specifically *through* its flat-chord-plus-single-normalize step and not through RMS's or GEO's geodesic `slerp` — unlike $\theta_e$, already screened, which is a property of the icosahedron itself rather than of SSB's particular discretization. An unnormalized icosahedron vertex at coordinates $(0, 1, \varphi)$ has Euclidean norm $s = \sqrt{1+\varphi^2} = \sqrt{\varphi+2}$ (using $\varphi^2 = \varphi + 1$); SSB's radial projection step divides by exactly this kind of unnormalized-chord length at every subdivision level, so $s$ is, in this sense, SSB's own signature constant in a way $\theta_e$ is not.

Building this basis correctly took two false starts, both caught by an internal PSLQ check (screening the basis alone, without $\kappa_\infty^{\mathrm{SSB}}$, before trusting any result against it) rather than by inspection. First, $\arctan(2)$ was initially included alongside $\arctan(1/\varphi)$ as two independent transcendental candidates; the internal check found $2\arctan(1/\varphi) - \arctan(2) = 0$ exactly (the classical tan-double-angle identity at $\tan\theta = 1/\varphi$: $\tan 2\theta = 2\tan\theta/(1-\tan^2\theta) = 2$ when $1 - \tan^2\theta = \tan\theta$, which holds here since $1 - 1/\varphi^2 = 1 - (\varphi + 2 - \varphi^2)/\varphi^2$ reduces to $1/\varphi$ using $\varphi^2 = \varphi + 1$), so only $\arctan(1/\varphi)$ was kept. Second, $\{s,\ 1/s,\ \varphi/s,\ \varphi s\}$ together were found to span only a two-dimensional space over $\mathbb{Q}$, not four independent directions, since $s^2 = \varphi + 2$ lets $1/(\varphi+2)$ itself be written as $3/5 - \varphi/5$ directly from $\varphi^2 = \varphi + 1$, making $1/s$ and $\varphi/s$ each a $\mathbb{Q}$-linear combination of $s$ and $\varphi s$ after clearing the denominator; the corrected basis keeps only $\{s, \varphi s\}$ from this family, dropping $1/s$ and $\varphi/s$, following the same in-family-redundancy discipline section 25.7.46 already established.

**Key Result**

*[A third family screened, and also empty] The corrected, internally-verified-independent twelve-term basis $\{1, \varphi, s, \varphi s, a_1, \pi,$ and each of $s, \varphi s, a_1$ divided by $\pi$ or $\pi^2\}$ (where $a_1 = \arctan(1/\varphi)$), screened against $\kappa_\infty^{\mathrm{SSB}}$ recomputed independently at $n = 10^{80}$, 220 working digits (agreeing with section 25.7.46's own recomputation to every digit printed), returns no relation, and the same basis screened against itself alone — confirming, this time, genuine $\mathbb{Q}$-independence rather than a hidden collision — also returns none.*

A quick supplementary check with `mpmath.identify` (a different, non-PSLQ search heuristic) against $\{\pi, \varphi\}$, against $\pi$ alone, and against $\{\pi, \sqrt{5}\}$ at 60 working digits agrees,

returning no candidate in any of the three cases — a weaker but independent corroboration, not a fourth basis in its own right.

**Todo — verify primary source**

Three independently-motivated families — $M_\infty^{(\text{area})}$'s own $\{\sqrt{3}, \sqrt{5}, \sqrt{15}, \pi\}$ ingredients (section 25.7.45), the icosahedron's defining angle and the golden ratio directly (section 25.7.46), and now SSB's own chord-normalization constant (section 25.7.47) — have each been screened at 220 digits and each come back empty. This narrows what a closed form for $\kappa_\infty^{\text{SSB}}$, if one exists, can look like, without narrowing the value itself: it is not a low-degree combination of any of the three most geometrically motivated ingredient sets this volume has been able to name. $\arccos(1/3)$, the digamma and elliptic-integral candidates the preceding todobox also named, and the unfitted Richardson-extrapolation question both remain exactly as open as before this section; nothing here privileges one over the other as the more promising next step.

### 25.7.48 A third convention: centroid-fan triangulation

Section 25.7.2's closing disclaimerbox, in earlier drafts of this chapter, named centroid-fan triangulation as the natural next convention to test and left it untested. This subsection tests it.

**Definition 25.60** (The centroid-fan convention)**.** For a Goldberg vertex $v$ and an incident face $F$ with vertex cycle $p_1, \ldots, p_m \ni v = p_i$, let $C = \frac{1}{m}\sum_j p_j$ be the ordinary Euclidean centroid of $F$'s vertices (the plain mean position in $\mathbb{R}^3$, *not* renormalized back onto the sphere). Decompose $F$ into the $m$ triangles $(C, p_j, p_{j+1})$; each is automatically planar, since any three points are. Define $F$'s contribution to the angle at $v$ as the sum of the two angles $v$ makes with $C$ within its two incident fan triangles, $\angle(p_{i-1}, v, C) + \angle(C, v, p_{i+1})$.

This is a different construction in kind from both earlier conventions — the raw convention (proposition 25.9) uses only $v$'s two edges within $F$ and ignores the rest of $F$ entirely, while the plane-projected convention (definition 25.10) flattens the whole of $F$ at once by a single best-fit plane — but it is exactly the ordinary interior angle whenever $F$ happens to be planar and convex, since the two fan triangles' angles at $v$ then sum to the true corner angle. It is a standard device in discrete differential geometry for assigning per-corner angles to a non-planar polygonal face (fan-triangulating from an interior point and summing), used for instance in cotangent-weight constructions on general polygonal meshes, and this chapter uses it here in that same standard sense rather than inventing it for the occasion.

**Proposition 25.61** (Computed, at the same three frequencies as table 25.1)**.** *Under definition 25.60's convention, the pentagon-corner defect is* $9.099926^\circ$ *at* $n = 2$, $3.684851^\circ$ *at* $n = 3$, *and* $1.908323^\circ$ *at* $n = 4$, *uniform across all sixty pentagon-corner vertices at each frequency to within* $1.6 \times 10^{-13}$ *degrees (the same order as the two earlier conventions' machine-precision uniformity). The total defect summed over* all *vertices of* $\mathcal{P}_n$ *is* $720.532395^\circ$ *at* $n = 2$, $720.222282^\circ$ *at* $n = 3$, *and* $720.175922^\circ$ *at* $n = 4$*: closer to* $720^\circ$ *as* $n$ *grows, but not equal to it at any of the three frequencies tested, unlike proposition 25.11's exact total under the plane-projected convention.*

**Key Result**

*[Pointwise, the third convention changes almost nothing; globally, it changes something real] Proposition 25.61's pentagon-corner values sit between the raw and plane-projected values at every frequency tested (*$9.105861^\circ > 9.099926^\circ > 9.098665^\circ$ *at* $n = 2$*, and the same ordering at* $n = 3, 4$*), and closer to the plane-projected value than the raw value is: at* $n = 2$ *the centroid-fan/plane-projected gap is* $0.001261^\circ$ *(about* $0.014\%$ *relative), smaller than section 25.7.2's already-small raw/plane-projected gap of* $0.007196^\circ$*. Pointwise, then, the disclaimerbox's expectation was right: a third reasonable convention does not disagree by much more, if anything it agrees somewhat better. But the total-sum check tells a different story from either earlier convention: proposition 25.7's* $720^\circ$ *identity, exact for the raw-angle convention's own idealized flat-faced hypothesis and verified exact for the plane-projected convention by proposition 25.11, is* not *matched by the centroid-fan convention at any of* $n = 2, 3, 4$*, and the mismatch — while shrinking with* $n$ *— does not appear to vanish at a rate that would put it within floating-point roundoff at the frequencies tested. Convention-dependence, in other words, is smaller than section 25.7.1's original todobox feared when measured pointwise at the sixty pentagon-corner vertices specifically, but is not uniformly negligible once the whole-polyhedron bookkeeping is checked, and the two checks can disagree about how much a convention "matters."*

*Why the total fails to close: an explanation, not merely an observation.* Proposition 25.7's proof requires an honestly flat-faced polyhedron, one in which every face, taken as a single planar polygon, contributes interior angles that literally sum the way section 25.2's cone-point construction assumes. Definition 25.10 manufactures exactly this by replacing each face with its own flattened image before measuring any angle, so proposition 25.11's exactness is inherited directly from proposition 25.7. Definition 25.60 does not: it decomposes each non-planar face into flat triangles meeting at an auxiliary point $C$ that is not itself a vertex of $\mathcal{P}_n$ and is not included anywhere in the vertex sum $\sum_v \delta(v)$. Whatever angular defect sits at $C$ under this fan decomposition — $C$ is, in general, itself a cone point of the fanned construction, with its own nonzero $2\pi$-minus-angle-sum defect whenever $F$ is non-planar — is simply omitted from the total rather than cancelling against anything, which is exactly the kind of bookkeeping gap that would produce a total systematically different from $720^\circ$ without any arithmetic error in the per-vertex computation itself. This is consistent with the sign and rough magnitude of what was found (a positive excess at every tested $n$, shrinking as the hexagons become flatter with refinement) without this chapter deriving the excess's exact value from $C$'s own implicit defect, which is a further computation this chapter has not carried out. □

### 25.7.49 Closing the gap: the face centroids, counted as their own cone points

Section 25.7.48's proof gave a qualitative mechanism for the centroid-fan convention's departure from $720^\circ$: each face's fan centroid $C$ is itself a cone point of the fan decomposition, and its defect is omitted from the vertex sum. This subsection makes that mechanism exact, by computing every such $C$'s own defect directly and checking the extended total.

**Definition 25.62** (Face-centroid defect)**.** For a Goldberg face $F$ with cyclically ordered vertices $p_1, \ldots, p_m$ and centroid $C = \frac{1}{m}\sum_j p_j$ (definition 25.60), define $C$'s own angular defect, exactly as at any cone point, as $\delta(C) := 2\pi - \sum_{i=1}^{m} \angle(p_i, C, p_{i+1})$, the angle sum being taken over the $m$ flat fan triangles meeting at $C$.

**Proposition 25.63** (The extended total is exactly $720^\circ$)**.** *For each of* $n = 2, 3, 4$*,* $\sum_v \delta_{\text{fan}}(v) + \sum_F \delta(C_F) = 720.000000^\circ$ *to within* $1 \times 10^{-6}$ *degrees, where the first sum is proposition 25.61's Goldberg-vertex total under the centroid-fan convention and the second is the sum of definition 25.62's defect over every one of* $\mathcal{P}_n$*'s faces.*

*Proof.* The fan decomposition of definition 25.60, applied to every face of $\mathcal{P}_n$ at once, replaces $\mathcal{P}_n$ by a strictly finer polyhedral surface: every original face is cut into $m$ flat triangles about its own new vertex $C_F$, every original Goldberg vertex $v$ retains its position, and the result is again a convex-faced (each face now a literal triangle, hence trivially planar) polyhedron combinatorially equivalent to a sphere, since the fan subdivision does not change the surface's topology. Proposition 25.7's hypothesis — a genuinely flat-faced polyhedron, combinatorially a sphere — is therefore met by this finer surface without qualification (unlike $\mathcal{P}_n$ itself under definition 25.60's convention, which is not literally flat-faced), and that proposition's proof, general in the polyhedron and requiring only proposition 25.1's angle-independent holonomy argument, applies to it verbatim: the sum of angular defects over *all* of the fine surface's vertices — the original Goldberg vertices, each now carrying exactly the centroid-fan angle sum of definition 25.60 at its corners, plus the new face-centroid vertices $C_F$, each carrying definition 25.62's defect — equals $4\pi = 720°$ exactly. The two computed sums proposition 25.61 and this subsection's own direct calculation of $\sum_F \delta(C_F)$ are independent numerical evaluations of the two halves of this one identity, and their agreement with 720° to six decimal places is a computational confirmation of the argument just given, not a separate coincidence. □

**Key Result**

*[The missing curvature sits at the hexagon centroids, and it is negative] At every frequency tested, the twelve pentagon-face centroids carry defect $\delta(C) = 0°$ exactly (to within $5 \times 10^{-14}$ degrees, machine precision), a direct consequence of section 25.7.1's exact planarity of every pentagon face: a planar $m$-gon's fan-triangulated angle sum at its own centroid is exactly $2\pi$ by elementary plane geometry, regardless of the polygon's specific shape, so a face with zero defect. The entire shortfall is carried by the hexagon-face centroids, and it is* negative*: at $n = 2$, all thirty hexagon centroids carry the identical defect $-0.017746°$ (uniform by the same twelve-fold symmetry that made section 25.7.1's planarity deviation uniform), summing to $-0.532395°$, exactly cancelling proposition 25.61's reported Goldberg-vertex excess of $+0.532395°$ at that frequency. At $n = 3, 4$, where the hexagons are no longer all symmetry-equivalent, the individual centroid defects vary (standard deviation $0.0016°$ and $0.0011°$ respectively, all values negative or, in a few cases at $n = 4$, indistinguishable from zero to the precision reported) but sum to exactly the negative of the corresponding vertex-sum excess in every case. A negative angular defect at $C$ means the fan triangles' angles at $C$ sum to* more *than $2\pi$: geometrically, $C$ sits at a saddle-shaped point of the fanned surface rather than a convex one, consistent with a non-planar hexagon (section 25.7.1) whose own centroid, computed as a plain Euclidean mean rather than fitted to any particular flattening, need not lie at a locally convex point of the fan it anchors.*

**Epistemic Notice**

Proposition 25.63 and section 25.7.49 close the todobox left open earlier in this section: the centroid-fan convention's departure from 720° is not an unexplained residual but the exact, now-computed defect of the sixty (at $n = 2$; more at higher $n$) auxiliary points this chapter's own construction introduces and had not, until this subsection, thought to count. This does not change section 25.7.48's substantive conclusion — the centroid-fan convention remains, as a convention for the sixty pentagon-corner *Goldberg*-vertex defects specifically, in close pointwise agreement with the other two — but it does mean the earlier framing of a convention that "does not restore exactness" was, in a precise sense, incomplete: exactness was restored all along, at a set of points this chapter had not yet included in the count.

**Epistemic Notice**

Two further scope limits on proposition 25.9 deserve stating plainly. First, the specific numbers $9.105861°$ and $121.4471°$ are computed for $\mathcal{P}_2$ (frequency $n = 2$) built by one specific, standard but not unique construction of class-I subdivision (edge-midpoint splitting with renormalization at a single level, rather than, say, a direct equal-arc geodesic grid on each icosahedral face); a different legitimate class-I construction at the same frequency, or the same construction at a different frequency $n \geq 2$, would in general return different specific numbers, though section 25.7.1's exact-$108°$ result and section 25.7.1's qualitative non-planarity should persist across constructions by the same symmetry argument. Second, section 25.7.2 above has since extended this computation to $n = 3, 4$ and to a second convention, so the "computes only $\mathcal{P}_2$" limitation this paragraph originally stated is superseded; what remains genuinely open is stated in section 25.7.2's own todobox and closing disclaimerbox, not repeated again here.

**Epistemic Notice**

Collecting this section's results in full: convergence of the underlying surfaces to $S^2$ is fully established (proposition 25.6); the total-curvature identity is exact at every finite stage of refinement for a genuinely flat-faced polyhedron and therefore requires no limiting argument at all, contrary to how the question is often informally posed, and this exactness has now been verified computationally at three frequencies under a convention that literally flattens every face (propositions 25.7 and 25.11); the number and location of the field's singularities is *not* fixed at twelve as refinement proceeds — it is twenty for the dodecahedron and sixty, at a fixed twelve-fold-symmetric arrangement, for every finer class-I subdivision (proposition 25.8); the pentagon-corner contribution to each of those sixty defects is exactly 108°, provably so by symmetry rather than merely observed (section 25.7.1); the full defect at each of those sixty vertices has been computed explicitly at three frequencies and under three conventions — a third, centroid-fan triangulation, added last (section 25.7.48) — that agree pointwise with each other far better than any of them agrees with the earlier idealized estimate, and whose apparent departure from an exact 720° total was traced to, and exactly accounted for by, the negative defect carried at each non-planar hexagon's own fan centroid, a set of auxiliary points the convention introduces but the vertex sum alone does not count (sections 25.7.1, 25.7.2, 25.7.48 and 25.7.49 and propositions 25.9 and 25.63); and, most unexpectedly, the curvature this construction actually carries does not stay concentrated at the twelve original pentagons as refinement proceeds, but spreads increasingly onto the hexagon-only vertices instead (section 25.7.2). None of this identifies a specific smooth vector field on $S^2$, nor an explicit sense in which the discrete singularities "become" a specific smooth field's zero set as $n \to \infty$, nor a proved asymptotic rate for any of the quantities in table 25.1; a least-squares fit over $n = 10$ to 40 narrows the rate to $-2.079 \pm 0.007$, closer to the conjectured $-2$ than earlier pairwise slopes suggested but still not equal to it within its own confidence interval (proposition 25.12 and section 25.7.4), and extending the fit one frequency further, to $n = 50$, narrows it again to $-2.049 \pm 0.005$ over $n = 20$ to 50, still short of $-2$ but monotonically closer at every successive restriction tested so far (proposition 25.13 and section 25.7.5). A further, and new, finding bears directly on the vector-field question itself: section 25.7.6's three shape statistics ($\mathrm{CV}(n)$, $M(n)$, $R(n)$, table 25.5) each move only slightly between $n = 20$ and $n = 50$, well short of the values exact uniformity would require, suggesting the defect field settles toward a fixed relative profile rather than flattening out entirely as $n \to \infty$ (section 25.7.6) — a reading sections 25.7.31 and 25.7.32's later extensions, first to $n = 100$ and then to $n = 120, 150, 200$, continue to support for $\mathrm{CV}(n)$ and $R(n)$, though $M(n)$'s own apparently exact plateau across five frequencies ($n = 50$ to 100) turns out not to be exact after all: it moves for the first time, from 1.206 to 1.207, once the range is pushed to $n = 200$ (section 25.7.32), a small but genuine correction to how firmly this chapter's own earlier drafts stated that particular plateau. If that reading holds, the original question of an explicit smooth vector field whose isolated zeros the discrete singularities converge to is not merely open but may be asking for the wrong kind of limiting object; the right one may be a continuous, non-constant curvature-density profile on $S^2$, not a finite zero set. All of this remains genuinely open — no closed form for the profile is proposed, and no proof that the shape statistics actually converge (rather than continuing to drift slowly) is given — but it is now open in a considerably sharper and partly quantified form than section 18.3 originally posed it in, and sharper again than earlier drafts of this chapter left it.

# Chapter 26

# The Closing Taxonomy: Four Axes, One Proved, Three Organizing

## 26.1 What this chapter can, and cannot, claim

Chapter 4's closing disclaimerbox anticipated a "two-axis argument" that would place every projection treated anywhere in this series into a single exhaustive grid, and flagged explicitly that the argument was not made there. Part III then added a third axis (discreteness, chapter 18) to the two anticipated (family, section 4.4; deployment, chapter 16), and, in its two final chapters, supplied the evidence for a fourth (continuity, chapters 19 and 20). This chapter makes the full argument, and its first task is to be precise about what kind of argument each of the four axes actually admits, because the four are not epistemically alike, and blurring that difference would be exactly the kind of overclaiming this series has tried to avoid since *No Map Is Best*'s first volume.

**Definition 26.1** (The four axes)**.** Every named construction this volume or its companions have treated is classified along four independent axes:

- **Family** (definition 4.4): conformal, equal-area, or aphylactic/compromise, determined pointwise (or cellwise, for a discrete construction, per section 18.2's extension) by the local scale pair $(a, b)$.

- **Deployment**: *global, single instance* (one continuous formula, or one discrete tiling, intended to represent the whole sphere at once, as with Mercator or the ISEA3H DGGS) versus *regional, re-centred* (the same formula reapplied with a shifted centre or standard line for each of several zones, as with Gauss–Krüger/UTM).

- **Domain**: continuous (a smooth map $f : U \to \mathbb{R}^2$ as in chapter 3's framework throughout Parts I–II) versus discrete (a finite cell partition as in chapter 18).

- **Continuity**: *contiguous* (the image of the sphere under the construction is a single connected region of the plane, as with every case treated before chapter 20) versus *interrupted* (the domain is cut, by design, into separate gores or lobes arranged non-contiguously, as with the Goode homolosine of chapter 20).

**Key Result**

*Only the family axis is exhaustive and exclusive as a consequence of a proved theorem: corollary 4.5 derives the trichotomy directly from theorems 3.1 and 4.1, so that no projection — named or hypothetical, treated in this series or not — can fall outside the three families or inside two of them at once, for reasons traceable all the way back to chapter 2's Theorema Egregium. The other three axes are exhaustive only in the weaker, definitional sense that the two options named in each case are constructed to be complementary by how this chapter has just defined them; nothing analogous to theorem 4.1 has been proved to force every conceivable cartographic construction into exactly one deployment strategy, exactly one domain type, or exactly one continuity class. The continuity axis in particular is the newest and least tested of the four: section 26.2 below adopts it on the strength of chapters 19 and 20's combined two-data-point argument, not a proof of orthogonality, and this chapter repeats that qualification rather than letting the axis's formal promotion in definition 26.1 above obscure it. The four-axis grid built in this chapter is therefore exhaustive in a mixed sense: provably so along one axis, definitionally so (with flagged exceptions, sharpest for the newest axis) along the other three.*

## 26.2 The full grid, populated

Definition 26.1's four axes generate $3 \times 2 \times 2 \times 2 = 24$ cells. Table 26.1 places the eight historical cases this volume's Part III has now assembled, together with the one non-historical construction chapter 23 added to close the grid's last open cell, within this grid, extending table 19.4 one last time by making the axis structure explicit as columns rather than as an evolving table format, and adding the Continuity column definition 26.1 just introduced.

| Construction | Family | Deployment | Domain | Continuity |
|---|---|---|---|---|
| Mercator, 1569 | Conformal | Global, single instance | Continuous | Contiguous |
| Lambert azimuthal, 1772 | Equal-area | Global, single instance | Continuous | Contiguous |
| Gauss/Krüger, 1822/1912 | Conformal | Regional, re-centred | Continuous | Contiguous |
| ISEA3H DGGS, 2003 | Equal-area | Global, single instance[1] | Discrete | Contiguous |
| Robinson, 1963/1974 | Aphylactic | Global, single instance | Continuous | Contiguous |
| Goode homolosine, 1923/1925 | Equal-area | Global, single instance | Continuous | Interrupted |
| Fuller/Dymaxion, 1943/1954 | Aphylactic | Global, single instance | Continuous | Interrupted |
| Peirce quincuncial, 1879 | Conformal | Global, single instance | Continuous | Contiguous |
| Stereographic-icosahedral (this volume) | Conformal | Global, single instance | Continuous | Interrupted |

Table 26.1: The complete four-axis placement of Part III's eight historical cases and the one non-historical construction chapter 23 adds. Of the twenty-four possible cells, eight are now populated (seven by a named historical case study, one — conformal-and-interrupted — by chapter 23's own construction rather than a historical map; Peirce's case, chapter 22, lands in a cell Mercator and Gauss/Krüger already occupy, rather than a new one), one of them (Fuller/Dymaxion, chapter 21) twice-classified along the family axis with the honest split section 21.4 recorded: closed-form proof for this volume's own gnomonic proxy, cited secondary evidence for the historical construction itself. The remaining sixteen cells are not thereby shown to be empty of real projections, only unrepresented by a dedicated case-study chapter here. Family and continuity are independently represented across all three family values (section 21.9); the fourth combination, conformal-and-interrupted, was examined directly by chapter 22 and found unoccupied by that chapter's own construction, for a reason section 22.10 argues is structural to Peirce's particular route to conformality rather than to conformality and interruption in general — a distinction chapter 23 then tests by building conformality a different way, and finds the cell is not, after all, structurally empty.

**Epistemic Notice**

Sixteen of the twenty-four cells in table 26.1 have no case-study chapter in this volume. Of the four gaps flagged across this volume's drafting — the aphylactic/compromise gap section 18.4 first raised, the interruption gap section 26.3 raised next, the specific aphylactic/interrupted cell section 20.5 and this section's earlier drafts both flagged as the narrowest remaining one, and the conformal-and-interrupted cell named once the first three were closed (this box's own earlier draft, and section 21.9's closing paragraph) — chapter 19 closed the first, chapter 20 closed the second, chapter 21 closed the third, and the fourth is now closed as well, though not by a historical case study. Chapter 22 examined the fourth cell directly using Peirce's quincuncial projection and did *not* close it: Peirce's construction, the natural historical candidate, turns out to be contiguous rather than interrupted, and section 22.10 argued this reflects a genuine tension between Peirce's own route to conformality (a single global conformal map, by Riemann's mapping theorem, which produces one connected image by its own construction) and the route to interruption chapters 20 and 21 use (cutting a net and rearranging pieces with gaps). Chapter 23 then tested whether that tension was intrinsic to conformal-and-interrupted constructions generally, or specific to Peirce's particular route, by building conformality a different way — stereographic projection applied gore by gore to the same icosahedral net chapter 21 used, rather than one global elliptic-function map — and found the cell populated after all: conformal at every point of every gore, by proposition 23.2, and interrupted by the same combinatorial cutting chapters 20 and 21 used. Section 21.9 is careful to note what the three historically-grounded cases that closed their cells establish: family and continuity have been observed to vary independently across three cases in detail (Robinson, Goode, Fuller), consistent with definition 26.1's orthogonal-axis reading, but three data points remain three data points, not a proof that no subtler coupling exists among cells this volume examines only through chapter 23's own constructed witness rather than a named historical map. Chapter 23's own closing disclaimerbox is explicit that this closure is narrower than the other three: it shows the cell is not structurally empty, not that every historically named conformal projection could have been interrupted, and Peirce's own resistance to interruption is left exactly as section 22.10 found it, sidestepped rather than explained away.

## 26.3 The edge case that motivated the fourth axis

Before chapters 19 and 20, this section recorded an open question rather than a resolution, and it is worth preserving that record rather than rewriting history now that definition 26.1 has adopted an answer. This volume's original three-axis definition, in the version this section first stated, defined deployment as a dichotomy between one global instance and several regionally re-centred instances of the same formula, and the interrupted projection — a single formula applied once, but with its domain cut along chosen meridians into separate, non-contiguously arranged gores or lobes, as in the Goode homolosine — fit neither value cleanly: not "regional, re-centred," because it is one formula applied once, not several instances centred differently; not straightforwardly "global, single instance" either, because interruption's whole point is that no single contiguous region of the plane represents the whole sphere continuously.

[1]A DGGS's hierarchy of resolutions could arguably be read as a form of "re-centring," each resolution level a distinct discrete grid; this chapter classifies deployment by whether the base construction covers the whole sphere in a single pass at a given resolution, which the ISEA3H (and every DGGS surveyed by Sahr, White & Kimerling) does, and treats multi-resolution hierarchy as a separate structural feature not captured by definition 26.1's four named axes at all — a fifth possible axis this chapter does not attempt to add, flagged explicitly rather than silently folded into deployment.

The question this box posed — three-valued deployment axis, or an independent fourth axis — is answered and adopted: definition 26.1 takes the fourth-axis reading, on the strength of chapter 19 and chapter 20's two data points showing family and continuity varying independently. Chapter 21 then supplied the third data point this box's own text anticipated wanting, and chapter 23 (added later still) supplied a fourth, populating the conformal-and-interrupted cell itself. Four cases now support the orthogonal-axis reading rather than two, though section 21.9's and chapter 23's own closing paragraphs are both correct to keep calling this strong support rather than a proof that no subtler coupling exists among cells this volume has not examined; that residual caveat is a standing feature of an empirical taxonomic claim, not an unresolved todo item, and is carried forward explicitly in table 26.1's own caption rather than left in this box.

## 26.4 How the taxonomy relates to this volume's four impossibility theorems

Chapter 1 organized this whole volume around four impossibility claims, listed as items (i)–(iv) in section 1.2. It is worth closing by making explicit how definition 26.1's taxonomy relates to each, since the taxonomy's purpose was never simply to catalogue named projections for their own sake, but to organize them as evidence of, and responses to, those four theorems specifically.

- Item (i), theorem 4.1, is the taxonomy's family axis directly: corollary 4.5's trichotomy *is* the family axis, and every entry in table 26.1's Family column is a witness to the theorem's content, not merely a label attached after the fact — now including Robinson, verified by computation rather than assumed, and Goode, verified by the exact closed-form proof proposition 20.1 gave its equatorial half.

- Item (ii), theorem 5.10 and theorem 5.2, does not appear as a column of table 26.1 at all, because it is a statement about the *magnitude* of unavoidable distortion within a region, not about which family a projection belongs to; its relationship to the taxonomy is that it quantifies, for any cell of the grid populated by a continuous-domain construction, exactly how much of the family axis's qualitative impossibility must be paid for locally — the same role proposition 4.3's trade-off curve already played for the family axis specifically, generalized by Milnor's theorem to arbitrary homeomorphisms rather than only conformal/equal-area pairs.

- Item (iii), the hairy ball obstruction of chapter 6, is the reason section 18.3 could prove proposition 18.1: both are consequences of $\chi(S^2) = 2$, and the taxonomy's Domain axis (continuous versus discrete) is where that shared ancestry becomes visible. Section 18.3's disclaimerbox flagged this ancestry as suggestive but not yet formally connected; chapter 25 closes that gap, showing both results are the sum-of-indices identity of the Poincaré–Hopf theorem evaluated on two different presentations (discrete cone-metric, continuous smooth) of the same sphere, rather than two separately-proved facts that merely happen to share an invariant.

- Item (iv), the retroazimuthal compatibility result of chapter 7, is not represented by any row of table 26.1 at all, and this is worth flagging as a further, specific gap rather than passing over silently: retroazimuthal projections are centred constructions, closer in spirit to the Deployment axis's "global, single instance" category (they are typically used whole, not zoned), but this volume's Part III chose not to give the retroazimuthal case its own case-study chapter, on the grounds that *The Retroazimuthal Problem* already treats it at

full length as a companion volume; this chapter does not reverse that choice, only records that item (iv), alone among the volume's four central results, has no representative in table 26.1.

**Epistemic Notice**

The four-item correspondence above is offered as an organizing summary of relationships already established piecemeal across this volume, not as a new theorem. Each bullet point restates a connection this volume's earlier chapters already made explicitly (the family axis's identity with corollary 4.5; Milnor's theorem as the quantitative sharpening of the family axis's trade-off; the shared invariant behind items (iii) and proposition 18.1) or explicitly declines to close (item (iv)'s absence from the grid). Presenting them together here is meant to make the volume's overall shape visible at a glance, not to claim anything not already stated, with its own qualifications, in the chapter where it was first established.

## 26.5 A fifth axis proposed, and four chapters outside the grid entirely

Five chapters were added to this volume after the draft this chapter otherwise describes: chapters 8 to 10 in Part II, and chapters 15 and 17 in Part III. None of them adds a row to table 26.1, and it is worth being precise about why, rather than leaving the omission to look like an oversight.

Chapter 8 is the one genuine candidate for a fifth axis: propositions 8.3 to 8.6 together describe, for any projection, the size of its exact-distance locus (empty, one point, two points, or provably-only-approximate beyond two), a classification orthogonal in principle to definition 26.1's existing four. Chapter 8 itself declines to fold this into table 26.1, on the grounds stated there: doing so properly means re-examining every entry in the table for its exact-distance locus specifically, which no chapter of this volume had done at the time that chapter was written (Mercator's and Gauss/Krüger's loci, for instance, were not stated anywhere in this volume, only assumed empty by family). Section 26.6 below now carries out that re-examination for all nine rows.

The remaining four chapters are not candidates for taxonomy rows at all, because they do not introduce named projections: chapters 9 and 10 generalize the curvature obstruction underlying the family axis itself to surfaces beyond the sphere, without constructing a new map of anything (a fifth chapter in this same non-row-adding family, chapter 11, was added later still, immediately after chapter 10; it examines an extrinsic realizability question rather than a projection and is excluded from table 26.1 for exactly the same reason); chapter 17 refines chapter 16's existing row (Gauss/Krüger, already in table 26.1) rather than adding one, showing the same family/deployment/domain/continuity classification survives the sphere-to-ellipsoid change unchanged (corollary 17.2); and chapter 15 examines a curve on the sphere, not a projection of it, so it has no natural row in a table of constructions at all. Table 26.1 accordingly stands as this chapter originally left it, sixteen of twenty-four cells unpopulated by anything this volume has itself examined as a dedicated case study — a count section 26.9 below narrows further, without adding a new case-study chapter for any of them, by surveying named constructions from outside this volume entirely — with the fifth axis now named, closed, and recorded in table 26.2 below rather than left as an open flag.

## 26.6 The fifth axis, closed: the exact-scale locus of every row in the table

Section 26.5's open item is closed here: every one of table 26.1's nine rows is examined directly for the size and shape of its exact-scale locus, rather than left, as the previous section put it, "assumed empty by family."

### 26.6.1 Two loci, not one, and why the distinction matters

Chapter 8 used "exact-distance locus" for a specific, strong property: a point $p$ such that $|f(x) - f(p)| = d_S(x, p)$ for *every* $x$ in the domain (proposition 8.4's global radial exactness). Chapters 3 and 8 separately used "distortion-free point" for a weaker, purely local property: a point $p$ with $a(p) = b(p) = 1$, i.e. the map is a local isometry at $p$ to first order (proposition 8.3). Every point satisfying the first property satisfies the second (radial exactness to every point forces isotropic unit scale at the center itself), but not conversely: proposition 26.3 below shows every point of Mercator's equator satisfies the second property, while chapter 8's own framework already implies no point of Mercator satisfies the first (a single projection cannot be simultaneously conformal, hence radially non-exact off any azimuthal center, and also exact-distance-preserving from a point on that center in every direction, except in the degenerate azimuthal-equidistant case itself). This section computes the second, weaker, always-well-defined locus — write $L(f) := \{p \in U : a(p) = b(p) = 1\}$ — for each of table 26.1's nine rows, and notes explicitly, case by case, whether the stronger global property also holds at any point of $L(f)$.

**Proposition 26.2** (Normalization dependence, stated once for all nine rows)**.** *For a conformal projection $f$ with local scale factor $\lambda(P)$ (so $a(P) = b(P) = \lambda(P)$ at every point where $f$ is conformal), the locus $L(f \cdot k_0) = \{P : \lambda(P) = 1/k_0\}$ for the uniformly rescaled map $k_0 f$ depends on the constant $k_0$: it can be empty, a single point, or a curve, according to whether $1/k_0$ lies outside, exactly at, or strictly below the infimum of $\lambda$ on $U$.*

*Proof.* Rescaling $f \mapsto k_0 f$ rescales $\lambda \mapsto k_0 \lambda$ uniformly, so $\{k_0\lambda = 1\} = \{\lambda = 1/k_0\}$; the three cases are the ordinary trichotomy of a level set of a continuous function against a target value compared to the function's infimum (below the infimum: empty; equal to the infimum, attained: possibly a single point or more, depending on where the infimum is attained; strictly above the infimum and below the supremum: generically a curve, by proposition 4.7 where the relevant regularity condition holds). □

This is not a new phenomenon specific to any one construction; it is exactly what UTM's secant form does to Gauss–Krüger's own exact locus (section 16.4: one curve becomes two, symmetric about the central meridian, once $k_0 = 0.9996$ replaces $k_0 = 1$), and it recurs below for Peirce, the one row where the historically used normalization constant has not been pinned down precisely enough in this volume to say which of proposition 26.2's three cases actually obtains.

### 26.6.2 The nine rows

**Proposition 26.3** (Mercator: a curve)**.** *$L(f)$ for the normal-aspect tangent Mercator projection is exactly the equator, $\{\varphi = 0\}$.*

*Proof.* Immediate from section 14.2's own keyresult: $\omega(P) \equiv 0$ everywhere (conformal family), and $s(P) = 1$ exactly along the equator and nowhere else, by corollary 4.5's forced failure of equivalence off that single standard parallel. Since $f$ is conformal, $s = \lambda^2$ with

$a = b = \lambda$, so $s = 1 \iff a = b = 1$ on the conformal locus, giving $L(f) = \{\varphi = 0\}$ exactly. □

Every point of this curve is, in the stronger sense too, exact-distance to every *other point of the same curve* (the equator is a great circle mapped isometrically to a straight line, proposition 8.6's remark on great-circle images applies), but not exact-distance to points off the equator; the equator is $L(f)$ in the local sense throughout, and additionally exact-distance in the strong sense restricted to itself, without being an azimuthal-equidistant center for the whole sphere.

**Proposition 26.4** (Lambert azimuthal equal-area: a single point)**.** *$L(f)$ for the Lambert azimuthal equal-area projection centered at $p_0$ is exactly $\{p_0\}$.*

*Proof.* $s(P) \equiv 1$ everywhere by section 14.4's keyresult (equal-area family), so $a(P)b(P) = 1$ identically; $a(P) = b(P)$ (conformality) then holds exactly where section 14.4's own argument places it, at the single center $p_0$, by corollary 4.5's exclusivity applied at every other point. Combined, $a = b = 1$ forces $a = b = 1$ exactly at $p_0$ and nowhere else. □

**Proposition 26.5** (Gauss–Krüger: a curve (or two))**.** *$L(f)$ for the tangent (basic) Gauss–Krüger projection is the central meridian; for the secant (UTM, $k_0 = 0.9996$) form it is the pair of curves section 16.4 already identifies, symmetric about the central meridian.*

*Proof.* Restated directly from section 16.4's own keyresult, itself an instance of proposition 26.2 with $k_0 = 1$ (tangent, single curve at the infimum) and $k_0 = 0.9996$ (secant, curve strictly above the infimum, split by proposition 4.7's regularity argument into two). □

**Proposition 26.6** (ISEA3H DGGS: not a locus of this kind)**.** *$L(f)$ as defined above — a subset of a continuous domain $U$ where the local Tissot ellipse degenerates to a unit circle — is not a well-posed question for the ISEA3H DGGS.*

**Epistemic Notice**

This is not a gap in this section's method but a structural fact about the construction itself, already flagged in chapter 18: the ISEA3H's equal-area property holds cell-by-cell (each discrete cell's area matches its spherical fraction exactly) rather than at a continuously varying point $P$ (section 18.3's own framing), so there is no continuous scale field $a(P), b(P)$ for $L(f)$ to be a level set of. A genuinely discrete analogue — for instance, which of the DGGS's own cell edges, if any, have exactly the spherical length of the corresponding geodesic arc — is a different, well-posed question this volume has not examined, and is recorded here as a new, narrower open item rather than answered by analogy with the continuous cases. Section 26.7 below now answers this item for the flat-faced polyhedral proxy this volume already uses for the ISEA3H family (proposition 26.13: no edge, at any resolution, ever matches its geodesic arc, by an elementary and fully general chord-versus-arc inequality), while explicitly leaving the corresponding question for Snyder's true, curved equal-area cell boundaries open in turn.

**Proposition 26.7** (Robinson: empty)**.** *$L(f) = \emptyset$ for the Robinson projection: exactly on the central meridian, by an argument that is genuinely exact; off it, by a closed-form monotonicity argument (proposition 26.8 below) that reduces the off-meridian case to the central-meridian one exactly, rather than by grid search alone.*

*Proof.* Section 19.3 and proposition 19.1 establish that, along the central meridian (where by the projection's own north–south symmetry the Jacobian is diagonal and $a, b$ coincide with the meridian and parallel scale factors $h, k$ directly, with no shear term to complicate the comparison), $k(\varphi) = 1$ only at $\varphi \approx 37.75°$–$37.82°$ and $\omega(\varphi) = 0$ (equivalently $h(\varphi) = k(\varphi)$) only at $\varphi \approx 32.85°$–$32.91°$ — two distinct latitudes, confirmed by two independent

interpolation methods. Since $a = b = 1$ requires both $h = k$ *and* $hk = 1$ simultaneously at the same point, and no latitude on the central meridian satisfies both, $L(f) \cap \{\lambda = 0\} = \emptyset$.

Off the central meridian, proposition 26.8 below shows $\omega(\varphi, \lambda) \geq \omega(\varphi, 0)$ for every $\lambda$ and every $\varphi \in (0°, 90°)$, so $\omega(\varphi, \lambda) = 0$ off the central meridian would require $\omega(\varphi, 0) = 0$ at the same latitude, already excluded by the preceding paragraph. This closes $L(f) \cap \{\lambda \neq 0\} = \emptyset$ directly from the central-meridian result, with no separate off-meridian search needed. `scripts/robinson_locus_search.py`'s earlier 64,619-point dense grid (from 0.5° to 89.5° latitude by 0° to 180° longitude) is retained below as independent numerical corroboration of proposition 26.8's closed-form conclusion, not as the argument's load-bearing step: it found zero violations of monotonicity at any of its 179 tested latitudes and a global minimum of 0.048526° at $\lambda = 0$, $\varphi = 32.5°$, exactly consistent with what the closed-form argument now proves outright. □

**Proposition 26.8** (Monotonicity in longitude, in closed form)**.** *For every $\varphi \in (0°, 90°)$, $\omega(\varphi, \lambda)$ is a non-decreasing function of $|\lambda|$, and strictly increasing wherever $X'(\varphi) \neq 0$.*

*Proof.* Equation (19.1)'s special pseudocylindrical form, $x = \mathrm{FXC}\, X(\varphi)\, \lambda$, $y = \mathrm{FYC}\, Y(\varphi)$, gives partial derivatives $x_\lambda = \mathrm{FXC}\, X(\varphi)$, $x_\varphi = \mathrm{FXC}\, X'(\varphi)\, \lambda$, $y_\lambda = 0$, $y_\varphi = \mathrm{FYC}\, Y'(\varphi)$ — all independent of $\lambda$ except $x_\varphi$, which is exactly linear in $\lambda$. In the orthonormal tangent frame ($u = \varphi$, $v = \cos\varphi \cdot \lambda$, so $\partial/\partial v = (1/\cos\varphi)\, \partial/\partial\lambda$), the Jacobian is

$$J(\varphi, \lambda) = \begin{pmatrix} A(\varphi) & B(\varphi)\, \lambda \\ 0 & D(\varphi) \end{pmatrix}, \qquad A = \frac{\mathrm{FXC}\, X(\varphi)}{\cos\varphi}, \quad B = \mathrm{FXC}\, X'(\varphi), \quad D = \mathrm{FYC}\, Y'(\varphi),$$

with $A, B, D$ all independent of $\lambda$. Writing $t = \lambda^2$, the two singular values $a, b$ of $J$ (the Tissot semi-axes) satisfy

$$a^2 + b^2 = \mathrm{tr}(J^\top J) = A^2 + D^2 + B^2 t, \qquad ab = |\det J| = |AD|,$$

the first non-decreasing in $t$ (since $B^2 \geq 0$), the second exactly constant in $t$. Since $\sin(\omega/2) = (a - b)/(a + b)$,

$$\sin^2(\omega/2) = \frac{(a-b)^2}{(a+b)^2} = 1 - \frac{4ab}{a^2 + b^2 + 2ab} = 1 - \frac{4|AD|}{A^2 + D^2 + B^2 t + 2|AD|},$$

which is non-decreasing in $t$ whenever $|AD| > 0$ (the denominator grows with $t$ while the numerator is fixed), strictly increasing whenever additionally $B \neq 0$. Since $\sin(\omega/2) \geq 0$ on $\omega \in [0, \pi]$, $\sin(\omega/2)$ itself, and hence $\omega$, inherits the same monotonicity in $t = \lambda^2$, i.e. in $|\lambda|$.

It remains to confirm $|AD| > 0$ throughout $\varphi \in (0°, 90°)$, i.e. that $X(\varphi) \neq 0$ and $Y'(\varphi) \neq 0$ there (the $\cos\varphi \neq 0$ factor in $A$ is automatic off the pole). `scripts/robinson_lambda_monotonicity_proof.py` checks both directly against eq. (19.1)'s own PROJ `x_robin.c` coefficients, and both are now closed-form exact rather than sampled. $Y'(\varphi)$'s sign: since $Y(\varphi)$ is cubic on each 5° segment, $Y'(\varphi)$ is a quadratic in the segment-local variable, whose real roots are exactly solvable in closed form; all eighteen segments have strictly positive $Y'$ at both endpoints and no real root of the quadratic falling inside $[0, 5]$, so $Y'(\varphi) > 0$ throughout each segment, hence throughout $(0°, 90°)$. $X(\varphi)$'s sign: since $X(\varphi)$ is itself cubic on each segment, $X(\varphi) = 0$ is solved directly by Cardano's closed-form formula (the trigonometric form when a segment's cubic has three real roots, Cardano's radicals when it has one), rather than by the dense sampling an earlier draft of this proof used; the resulting real roots are checked against $[0, 5]$ exactly, not approximately, and none falls in range on any of the eighteen segments, with the final segment reaching 0.5322 at $\varphi = 85°$ and not returning to 0 until the pole itself (excluded from the domain, as elsewhere in this chapter). This Cardano computation is cross-checked two

ways — against `numpy.roots` (companion-matrix eigenvalues, a numerically independent method) and, separately, against `scripts/robinson_X_sympy_crosscheck.py`'s sympy exact-rational Sturm-sequence real-root isolation, which treats the published coefficients as exact rationals and isolates real roots without ever evaluating the polynomial in floating point — and all three methods agree on every segment. □

**Epistemic Notice**

This closes exactly the item the previous draft of this disclaimerbox left open: a fully continuous proof of monotonicity in $\lambda$, not a denser grid. Two things are worth being precise about regarding what kind of proof this is. First, the monotonicity itself (proposition 26.8's main argument) is exact algebra with no numerical component at all, once $|AD| > 0$ is granted — it follows from $B^2 \geq 0$ and $|AD|$'s $\lambda$-independence alone, both immediate from eq. (19.1)'s specific pseudocylindrical form. Second, $|AD| > 0$ itself now rests on two checks of the *same* character, where an earlier draft's rested on two of different character: $Y'(\varphi) \neq 0$ is closed-form exact (a quadratic's roots, solved directly), and $X(\varphi) \neq 0$ is now closed-form exact as well (each segment's cubic solved directly by Cardano's formula, triple-cross-checked against `numpy.roots` and against sympy's exact-rational Sturm-sequence isolation), in place of the earlier draft's 1001-point-per-segment dense sampling. No residual sampling-based link remains in this proof: proposition 26.8's full chain, from the pseudocylindrical Jacobian form down to $X(\varphi) \neq 0$ and $Y'(\varphi) \neq 0$ on every one of the eighteen segments, is closed-form throughout.

**Proposition 26.9** (Goode homolosine: a cross in the sinusoidal band, plus two points per gore at the join)**.** *Within each gore's sinusoidal (equatorial) piece, $L(f)$ is the union of that gore's own central meridian and its segment of the equator, both restricted to $|\varphi| < 40°44'11.8''$. Within each gore's Mollweide (polar) piece, $L(f)$ is a single point on the central meridian, at the latitude where Mollweide's own parallel scale equals* 1 *— computed here at $\varphi \approx 40.7367°$, indistinguishable to four significant figures from Goode's own $40°44'11.8'' = 40.7366°$ join latitude.*

*Proof. Sinusoidal piece.* Writing the projection as $x = R\lambda\cos\varphi$, $y = R\varphi$ and passing to the orthonormal tangent frame $(u, v) = (R\lambda\cos\varphi, R\varphi)$-arclength coordinates as proposition 20.1's proof already sets up, the Jacobian $\partial(x, y)/\partial(u, v)$ has first column $(1, 0)$ identically (since $x_\lambda/(R\cos\varphi) = 1$ and $y_\lambda = 0$ exactly, both already established there) and second column $(-\lambda\sin\varphi, 1)$ (from $x_\varphi/R = -\lambda\sin\varphi$ and $y_\varphi/R = 1$). This matrix is orthogonal — the condition for $a = b = 1$ — exactly when its second column has unit norm and is orthogonal to the first, both equivalent to the single scalar condition $\lambda\sin\varphi = 0$, i.e. $\lambda = 0$ or $\varphi = 0$: the central meridian or the equator. `scripts/taxonomy_fifth_axis_check.py` confirms this by direct numerical singular-value decomposition of the Jacobian at eight sample points, on and off the predicted cross, agreeing with the closed-form criterion at every point tested (the two off-cross points give $(a, b) \approx (1.094, 0.914)$ and $(1.039, 0.963)$, visibly $\neq (1, 1)$; the on-cross points give $(1, 1)$ to six decimal places).

*Mollweide piece.* Analogously, with $y_\lambda = 0$ identically (already established in section 20.3), the Jacobian's off-diagonal terms both vanish exactly on the central meridian ($\lambda = 0$), where $a = b$ reduces to $h = k$; since Mollweide is equal-area everywhere (proposition 20.2, $hk \equiv 1$ identically), $h = k$ there forces $h = k = 1$ exactly where the parallel scale factor $k(\varphi) = \frac{2\sqrt{2}}{\pi}\cos\theta(\varphi)/\cos\varphi$ — section 20.3's own already-computed quantity — equals 1. Solving this numerically (bisection on the same $\theta(\varphi)$ relation section 20.3 already solves) gives $\varphi \approx 40.7367°$, matching section 20.4's independently derived join latitude ($40°44'11.8'' = 40.7366°$) to four significant figures — not a coincidence this volume is discovering for the first time, but the numerical confirmation of exactly the mechanism section 20.4's own text already states in words ("joined … exactly where each piece's own linear scale along the parallel equals 1"), now cross-checked against the full Tissot condition

rather than the parallel-scale factor alone. Off the central meridian, the same script's Jacobian SVD at the standard-parallel latitude gives $(a, b) \approx (1.055, 0.948)$ at $\lambda = 15°$ and $(1.153, 0.868)$ at $\lambda = 40°$, confirming the shear term the by-hand argument predicts breaks full isometry away from $\lambda = 0$ even along a parallel where the scalar parallel-scale factor alone equals 1. □

**Epistemic Notice**

The near-coincidence between Mollweide's own true-parallel-scale latitude and Goode's chosen join latitude is exact enough (four significant figures, using two independently computed quantities — section 20.4's own root-finding and this section's separate numerical solve of the same underlying $\theta(\varphi)$ relation) that it is very unlikely to be accidental, and is consistent with Goode's own stated design rationale as section 20.4 already reports it; this section does not have a primary-source confirmation that Goode chose the join *because* of this specific point-locus property (as opposed to only the scalar parallel-scale coincidence already documented), and states the numerical match without attributing that stronger intent to Goode.

**Proposition 26.10** (Fuller/Dymaxion (gnomonic proxy): twenty points)**.** *$L(f) = \emptyset$ pointwise except at the twenty face centroids $\{\widehat{n}_F\}_F$, where $a = b = 1$ trivially.*

*Proof.* Immediate from propositions 21.2 and 21.3: $h(c) = \sec^2 c$, $k(c) = \sec c$, both equal to 1 exactly at $c = 0$ (the tangent point) and strictly greater than 1, and unequal to each other, for every $c \in (0, \pi/2)$. Applied independently within each of the twenty faces (the construction is azimuthal per face, centered at that face's own centroid), this gives exactly one point per face. □

**Proposition 26.11** (Peirce quincuncial: empty, one point, or a closed curve, depending on a normalization this volume has not pinned down)**.** *Write the composed Peirce map as (stereographic projection from each point's antipode) ∘ (the disk-to-square elliptic map of proposition 22.4). Since both stages are conformal (away from the four singular points), the composed local scale factor is $\lambda_{\text{total}}(P) = |C| \cdot \sec^2(c(P)/2) \cdot |1 - z(P)^4|^{-1/2}$, where $c(P)$ is angular distance from the pole and $z(P)$ is the corresponding point in the unit disk. This function equals $|C|$ exactly at the pole ($c = 0$, $z = 0$) and increases monotonically (confirmed numerically along a representative radial path in `scripts/taxonomy_fifth_axis_check.py`, from $|C|$ at the pole to unboundedly large values approaching each of the four singular points, where $|1 - z^4|^{-1/2} \to \infty$). Consequently, by proposition 26.2: if $|C| > 1$, $L(f) = \emptyset$; if $|C| = 1$, $L(f) = \{\text{pole}\}$, a single point; if $|C| < 1$, $L(f)$ is a single closed curve encircling the pole, with the same four-fold symmetry as proposition 22.4's square.*

**Epistemic Notice**

Which of the three cases applies to Peirce's own historically published map, or to Pierpont's corrected version, was not settled at the point this proposition was first drafted: it depends on the specific numeric relationship between the physical size of the printed square and the sphere's radius $R$ in Peirce's own tables, tied to the same $K/2 = 1$ normalization detail section 22.2 had then only flagged as unresolved without a primary reading of Peirce's 1879 formula. What this proposition added at that stage was the qualitative structure — exactly one of three cases obtains, cleanly determined by a single constant, with monotonicity guaranteeing no more exotic locus shape (multiple disjoint curves, for instance) is possible — in place of no statement at all. Section 22.7 has since read Peirce's own printed tables directly and found his own $K/2 = 1$ normalization matches this volume's independently derived $|C| = 1/K$ convention exactly, which section 22.5 already shows places $|C| < 1$: the third case, a nonempty closed curve, is accordingly confirmed for Peirce's own historically published map specifically, not only for an independently motivated candidate convention.

**Proposition 26.12** (Stereographic-icosahedral: twenty points)**.** *$L(f) = \emptyset$ pointwise except at the twenty face tangent points $\{\widehat{n}_F\}_F$, exactly as for proposition 26.10, but for a different structural reason.*

*Proof.* By proposition 23.2, $\omega(P) \equiv 0$ on every gore's interior (the whole construction is conformal, unlike Fuller's proxy), so $a = b = \lambda(P) = \sec^2(c(P)/2)$ throughout, by proposition 22.2's own closed form applied per face. This equals 1 exactly at $c = 0$ and is strictly greater than 1 for every $c \in (0, \pi/2)$, giving one point per face — the same count as proposition 26.10, but here because area scale alone fails away from center on an otherwise-everywhere-conformal map, rather than because both angle and area fail together on a nowhere-conformal one. □

### 26.6.3 Summary

| Construction | $L(f)$ (local Tissot-exact locus) | Reference |
|---|---|---|
| Mercator, 1569 | Curve (the equator) | proposition 26.3 |
| Lambert azimuthal, 1772 | Point (the center) | proposition 26.4 |
| Gauss/Krüger, 1822/1912 | Curve (central meridian; two curves in UTM's secant form) | proposition 26.5 |
| ISEA3H DGGS, 2003 | Not applicable (discrete, cell-level equal-area only) | proposition 26.6 |
| Robinson, 1963/1974 | Empty (exact on central meridian; dense 64,619-point grid off it, zero exceptions) | proposition 26.7 |
| Goode homolosine, 1923/1925 | Cross (sinusoidal band) + one point per gore (Mollweide cap) | proposition 26.9 |
| Fuller/Dymaxion, 1943/1954 | Twenty points (face centroids, gnomonic proxy) | proposition 26.10 |
| Peirce quincuncial, 1879 | Nonempty closed curve encircling the pole (confirmed for Peirce's own historical convention) | proposition 26.11 |
| Stereographic-icosahedral (this volume) | Twenty points (face tangent points) | proposition 26.12 |

Table 26.2: The fifth axis, applied to all nine rows of table 26.1: the local Tissot-exact locus $L(f) = \{P : a(P) = b(P) = 1\}$, distinct from (and in every row here, a strict superset in the sense of section 26.6's opening discussion of) the stronger global exact-distance-from-a-point property chapter 8 characterizes for the azimuthal-type constructions specifically.

**Epistemic Notice**

This closes section 26.5's open item at the level it was posed — every named row now has a stated locus, not an assumption — but opened, at the time this section was drafted, two narrower items in its place, recorded here rather than left implicit. The first, whether Robinson's locus is empty everywhere or only on the central meridian, is now resolved at the standard of evidence this chapter uses throughout (exact on the central meridian, and a 64,619-point dense grid finding zero exceptions off it, proposition 26.7's own disclaimerbox), leaving only the narrower, fully explicit residual gap between adjacent grid nodes rather than the original all-of-$\mathbb{R}^2$ uncertainty. The second, which of proposition 26.11's three cases applies to the historically published Peirce/Pierpont formula specifically, was narrowed but left open at that stage, though a further pass (section 22.5) had by then already fixed $|C|$ under one specific, cross-checked convention distinct from Peirce's own historical tables — vertex-to-center distance 1, matching an independently normalized source outside this volume's prior bibliography — and found $|C| = 1/K \approx 0.762760 < 1$ under that convention, which proposition 26.11's own trichotomy places squarely in the third case, a nonempty closed curve (section 22.5). At that stage this was progress of the same sharpening-without-closing kind this section's other items record, one well-motivated candidate answered, the historical-convention question still open, not the closure the historical question itself needed. That closure has since arrived: section 22.7 reads Peirce's own printed tables directly and finds his historical maximum matches the vertex-unit convention exactly, so the case-three conclusion above is no longer contingent on an independently chosen candidate — it is Peirce's own historically published convention, confirmed against his own printed numbers. The Peirce row of table 26.2 is updated accordingly, and this item is closed in full, not merely narrowed. Both items were, in any case, always narrower and more tractable than the original "re-examine every entry" item, which is the sense in which this section is a closure rather than a restatement of the gap under a new name.

## 26.7 A discrete analogue for the ISEA3H family: do any cell edges match their geodesic arc?

Proposition 26.6's disclaimerbox left a new, narrower open item in place of the ill-posed continuous question it declined to ask of the ISEA3H DGGS: whether any of the DGGS's own cell edges, at any resolution, have exactly the spherical (geodesic) length of the corresponding great-circle arc between their endpoints. This section answers that question for the one object this volume can examine directly without further access to Snyder's own equal-area transfer formula: the flat-faced polyhedral proxy $\mathcal{P}_n$ that definition 25.5 already builds and section 25.7 uses throughout chapter 25 as a stand-in for the ISEA3H family's underlying icosahedral Goldberg structure — vertices of the Goldberg polyhedron $\mathrm{GP}(n, 0)$ placed on the unit sphere $S^2$ at their radial direction, joined by straight (planar, chordal) edges rather than by the curved cell boundaries Snyder's true equal-area transfer would produce.

**Proposition 26.13** (No proxy edge ever matches its geodesic arc)**.** *Let $P, Q$ be any two distinct points of $S^2$ that are not antipodal to each other, and let $\theta \in (0, \pi)$ be the angle between them (equivalently, the geodesic arc length between them, since $S^2$ has unit radius). Then the straight-line chord length $|P - Q|$ is strictly less than $\theta$. Consequently, for every $n \geq 1$ and every edge of $\mathcal{P}_n$, the flat edge's (chord) length is strictly less than the spherical (geodesic) arc length between its two endpoints: no edge, at any resolution, ever matches its geodesic arc exactly.*

*Proof.* Place $P, Q$ so that the angle between them, measured from the origin, is $\theta$. The chord length is $|P - Q| = 2\sin(\theta/2)$, an elementary consequence of the isosceles triangle $OPQ$ with $|OP| = |OQ| = 1$ and apex angle $\theta$ (drop the perpendicular bisector from $O$ to $PQ$; each half of the chord has length $\sin(\theta/2)$). The claim is therefore $2\sin(\theta/2) < \theta$ for $\theta \in (0, \pi)$, equivalently $\sin x < x$ for $x = \theta/2 \in (0, \pi/2)$. This is the standard strict inequality $\sin x < x$ for all $x > 0$ (immediate from the Taylor series $\sin x = x - x^3/3! + x^5/5! - \cdots$ with strictly negative remainder after the leading term for $x > 0$, or from the elementary unit-circle comparison of an arc to its chord), applied on the sub-interval $(0, \pi/2)$ where it is needed here. Since every edge of $\mathcal{P}_n$ joins two distinct vertices of the unit sphere at some angular separation $\theta_e \in (0, \pi)$ — distinct because $\mathcal{P}_n$ is a simple (non-degenerate) polyhedron with no repeated vertices, and strictly less than $\pi$ because no two vertices of $\mathcal{P}_n$, at any tested resolution or in general, are antipodal (the icosahedral Goldberg construction places all vertices on one hemisphere-scale neighborhood of each face, far short of the full angular range needed to reach an antipodal pair) — the general inequality applies edge by edge, giving $\mathrm{chord}_e < \mathrm{arc}_e$ for every edge $e$ of $\mathcal{P}_n$, at every $n$. □

**Epistemic Notice**

This proof is genuinely general: it does not depend on $n$, on the icosahedral construction specifically, or on any property of $\mathcal{P}_n$ beyond having straight edges between distinct, non-antipodal points of $S^2$ at all — the same argument applies verbatim to any flat-faced polyhedron inscribed in the unit sphere, DGGS-derived or not. It is included here as a full proof rather than only as a citation to a standard fact because section 26.5's own working method throughout this volume has been to prove elementary steps explicitly rather than wave at their standardness, and because the general inequality is precisely what closes the specific, narrower item proposition 26.6's disclaimerbox named — the proof is short exactly because the underlying fact is elementary, not because a step has been skipped.

`scripts/dggs_edge_chord_arc.py` confirms proposition 26.13 against $\mathcal{P}_n$'s actual computed edges (reusing the same icosahedron/subdivision/dual-face construction already validated in chapter 25, rather than a separately re-derived one), at eight resolutions from $n = 2$

to $n = 20$: across 26,220 edges checked in total, zero violations (no edge with chord $\geq$ arc) at any resolution. Table 26.3 records the relative gap (arc$-$chord)/arc at each resolution, which the proof's own small-angle expansion predicts should shrink like $\theta^2/24$ as edges get shorter with increasing $n$ (from $2\sin(\theta/2) = \theta - \theta^3/24 + O(\theta^5)$, giving $(\theta - 2\sin(\theta/2))/\theta \approx \theta^2/24$): the script confirms this rate directly, the ratio of the observed smallest-edge gap to the $\theta^2/24$ prediction converging to 1.0000 by $n = 10$ and holding there through $n = 20$.

| $n$ | Edges checked | Violations | Shortest edge arc (deg) | (arc − chord)/arc, shortest edge |
|---|---|---|---|---|
| 2 | 120 | 0 | 19.898° | $5.018 \times 10^{-3}$ |
| 3 | 270 | 0 | 12.582° | $2.008 \times 10^{-3}$ |
| 4 | 480 | 0 | 8.902° | $1.006 \times 10^{-3}$ |
| 5 | 750 | 0 | 6.818° | $5.899 \times 10^{-4}$ |
| 8 | 1920 | 0 | 3.957° | $1.987 \times 10^{-4}$ |
| 10 | 3000 | 0 | 3.083° | $1.206 \times 10^{-4}$ |
| 16 | 7680 | 0 | 1.850° | $4.346 \times 10^{-5}$ |
| 20 | 12000 | 0 | 1.460° | $2.707 \times 10^{-5}$ |

Table 26.3: Chord-versus-arc check for every edge of $\mathcal{P}_n$ at eight resolutions (26,220 edges total, zero violations). The relative gap shrinks toward zero as edges get shorter at higher $n$, exactly at the $\theta^2/24$ rate the proof's own Taylor expansion predicts, but — consistent with proposition 26.13 itself — never reaches zero at any finite $n$: the proxy's edges approach, but by proposition 26.13 can never attain, their geodesic length.

**Epistemic Notice**

This closes the item exactly as narrowly as it was posed for the flat-faced polyhedral proxy $\mathcal{P}_n$, and no further: $\mathcal{P}_n$ is a convenient stand-in this volume already uses for the ISEA3H family's underlying combinatorial structure, but it is not the ISEA3H DGGS itself, whose actual cells (per section 18.2) are bounded not by straight chords between Goldberg-polyhedron vertices but by curved boundaries produced by Snyder's 1992 equal-area transfer, a formula chapter 18 has already flagged, twice, as cited at the level of its existence and stated purpose rather than reproduced or verified in this volume. Proposition 26.13 says nothing directly about whether one of *those* curved boundaries could ever coincide with its geodesic arc; the proxy's flat chords are, if anything, the case most obviously biased toward inequality (a straight line is the unique shortest path between two points in the ambient $\mathbb{R}^3$, always at least as short as any curved path joining them, so chord $\leq$ arc is actually forced trivially for any flat-edged polyhedron on $S^2$, without needing proposition 26.13's $\sin x < x$ argument at all — what that argument earns is the *strict* inequality, ruling out the one case trivial subadditivity leaves open, namely equality). Snyder's own curved cell edges are subject to no such trivial bound either way, since they need not be geodesics or chords at all, and whether any of them ever matches its own geodesic arc exactly is recorded here, explicitly, as still open — a question this volume cannot close without the primary formula chapter 18's own todobox already lists as unread.

## 26.8 Snyder's actual curved boundaries: which, if any, are geodesics?

Chapter 18's own todobox notwithstanding, this section takes the question the previous disclaimerbox left open as far as it can be taken without a primary reading of Snyder's 1992 paper itself. The route in is a secondary reconstruction rather than the primary source: Brenton R. S. Recht's 2021 vector-form restatement of Snyder's equal-area transfer, itself checked there against Harrison, Mahdavi-Amiri & Samavati's independent 2012 analysis of the same

projection's inverse.[2] This is the same epistemic status as chapter 21's Gray reconstruction from posted source code: a specific, checkable restatement of a primary formula this volume has not itself read, used at face value and verified internally rather than against Snyder's own worked examples. The restated forward map passed its own defining boundary check directly (it sends the spherical triangle's three vertices to the three barycentric unit points, exactly); the same source's closed-form inverse formula did not reproduce one of the same three vertices correctly under this reading, a discrepancy this pass could not resolve from the blog post's abbreviated presentation alone, so every curved-boundary point used below comes instead from numerically inverting the validated forward map (root-finding, each solution re-checked against its own target to machine precision) rather than from the unverified closed-form inverse. Two further correctness pitfalls surfaced and were caught during this construction, both recorded here rather than silently fixed, in keeping with this volume's practice of showing its own errors: the vector identity used for the forward map's intermediate point conflates two different quantities under casual reading — a plain scalar triple product (used correctly only inside that one identity) and the spherical triangle's genuine excess area (needed everywhere else the formula calls for an area) — and treating the two as interchangeable silently breaks the area-additivity Snyder's own derivation depends on; the fix is to compute spherical excess area via the exact Van Oosterom–Strackee solid-angle formula rather than the triple product, verified directly against additivity before use.

**Proposition 26.14** (The apex median is exactly geodesic)**.** *Let $v_0, v_1, v_2$ be the vertices of one of Snyder's isosceles sub-triangles for the icosahedron (apex $v_0$ at the sub-triangle's own axis of symmetry, base $v_1, v_2$ two vertices of the icosahedron equidistant from $v_0$), and let $\mathcal{M}$ be the locus of points in the sub-triangle whose barycentric coordinates under Snyder's forward map satisfy $\beta_1 = \beta_2$. Then $\mathcal{M}$ is exactly the great-circle arc from $v_0$ to the geodesic midpoint of $v_1, v_2$: every point of $\mathcal{M}$ is a geodesic distance match, not merely an approximate one.*

*Proof.* Let $R$ be the reflection of $S^2$ through the plane spanned by $v_0$ and $v_1 + v_2$ (well-defined since $v_0, v_1, v_2$ are not collinear); $R$ fixes $v_0$ (equidistant from $v_1, v_2$ by construction, hence on the mirror plane) and swaps $v_1 \leftrightarrow v_2$. Two facts about Snyder's forward map, call it $F$, combine to place $\mathcal{M}$: first, relabeling $v_1 \leftrightarrow v_2$ in $F$'s own definition swaps $\beta_1 \leftrightarrow \beta_2$ in the output, for any point, since $F$'s intermediate point $\hat{p}$ (the intersection of the great circles through $v_0, v$ and through $v_1, v_2$) does not depend on which of $v_1, v_2$ is named first, while the area ratio defining $\beta_2$ becomes the area ratio defining $\beta_1$ under the swap by the additive splitting $A(v_0, v_1, \hat{p}) + A(v_0, \hat{p}, v_2) = A(v_0, v_1, v_2)$ (spherical excess area, exact, verified directly in section 26.8's companion script before use). Second, $F$ is intrinsic — built entirely from dot products, cross products, and spherical excess areas of the vectors involved, all quantities an isometry of $S^2$ carries to the corresponding quantities of the transformed vectors, with signed quantities changing sign under the orientation-reversing $R$ exactly as needed to cancel against the relabeling in the first fact. Composing the two: a point $P$ with barycentric $(\beta_0, \beta_1, \beta_2)$ under $(v_0, v_1, v_2)$ has $R(P)$ with barycentric $(\beta_0, \beta_2, \beta_1)$ under the same, unswapped labeling. For $P \in \mathcal{M}$ (where $\beta_1 = \beta_2$), this says $R(P)$ has the same barycentric triple as $P$; since $F$ is a bijection of the sub-triangle, $R(P) = P$, i.e. every point of $\mathcal{M}$ is fixed by $R$. The fixed-point set of a reflection of $S^2$, restricted to any connected piece of the sphere it meets, is exactly the great circle cut out by the reflection's mirror plane — here, the great circle through $v_0$ and the geodesic midpoint of $v_1, v_2$ (the one point on arc $v_1 v_2$ itself fixed by $R$, since $R$ is an isometry and isometries preserve geodesic midpoints

[2] B. R. S. Recht, "Snyder's equal-area projection," *The BRSR Blog*, 31 Aug 2021, consulted 6 Aug 2026; cross-checking Snyder (1992) against E. Harrison, A. Mahdavi-Amiri & F. Samavati, "Analysis of Inverse Snyder Optimizations," *Transactions on Computational Science XVI*, LNCS 7380, Springer, 2012.

regardless of what any particular map does to barycentric coordinates). □

**Key Result**

*[One family of ISEA3H grid lines is exactly geodesic — and the numbers confirm it to machine precision]* `scripts/dggs_snyder_curved_edges.py` *builds a genuine icosahedron (exact, golden-ratio vertex coordinates, no invented data), the isosceles sub-triangle Snyder's own construction specifies for one of its faces, and samples* 25 *points along* $\mathcal{M}$ *via the verified numerical inversion described above. Every sampled point satisfies* $|\hat{n} \cdot \hat{v}| < 3 \times 10^{-16}$ *against the great-circle plane proposition* 26.14 *names, where* $\hat{n}$ *is that plane's unit normal — floating-point-exact agreement, not merely small. Repeating the same check on an entirely independent icosahedron face gives the same result to the same precision, confirming this is not an artifact of one face's particular coordinates.*

Symmetry forces exactly this one family of grid lines onto geodesics — the apex medians of Snyder's own isosceles sub-triangles, which trace 60 great-circle arcs across the icosahedron (3 per face × 20 faces, radiating from each face's centroid toward its three edge midpoints). These are genuine ISEA3H grid lines (they separate a hexagon from its neighbour along the axis Snyder's construction itself singles out), so the answer to section 26.7's closing question is not simply "no": some curved boundaries do coincide with their geodesic arc, exactly, and provably so by the symmetry argument above rather than by numerical coincidence.

Every *other* tested boundary is a different matter. Three concrete, non-median aperture-3 grid chords — interior cell-boundary segments at barycentric grid points with denominators of three, chosen to avoid $\mathcal{M}$ and the sub-triangle's own outer edges (which are geodesic trivially, being spherical-triangle sides by construction) — were sampled the same way and checked against the unique great-circle plane through each chord's own two endpoints. All three show a clear, non-negligible departure from that plane: roughly 75, 211, and 1037 arc-seconds respectively at the sampled interior points, five to six orders of magnitude above the $\sim 10^{-15}$ residual of the numerical inversion itself, so the departures are a genuine geometric effect and not solver noise. On the WGS84 authalic sphere (radius 6,371,007.18 m, PROJ's own published constant for this exact purpose)[3] these translate to roughly 2.3, 6.5, and 32 kilometres of cross-track deviation from the geodesic — not a fine distinction near the limits of measurement, but a substantial, physically meaningful curvature. The same three-chord test repeated on the second, independent icosahedron face reproduces the identical deviation for the shape tested there, again to machine precision, ruling out a face-specific coincidence.

[3]PROJ contributors, "Icosahedral Snyder Equal Area," PROJ documentation, consulted 6 Aug 2026: the projection is defined only on a sphere, and PROJ's own recommendation for mapping WGS84 coordinates onto it is this authalic radius.

**Epistemic Notice**

What this closes, and what it does not. The apex-median result is a genuine closed-form fact, proved by symmetry and confirmed numerically rather than asserted from the numerics alone: an entire family of 60 real ISEA3H grid lines are geodesics, exactly. The three non-median chords tested establish, with a clear and quantified margin, that not every curved boundary is a geodesic — answering section 26.7's question in the negative for every configuration actually tested, robustly across two independent faces. What remains open is the fully general classification: whether some *other*, untested chord — at a different resolution, a different position within the sub-triangle, or arising from the pentagonal cells at the twelve icosahedron vertices rather than the hexagonal interior — might coincide with its geodesic by some further symmetry this pass has not identified. Proposition 26.14's proof technique (reflection symmetry forcing a fixed locus onto a great circle) is specific to the one reflection the isosceles construction actually admits; nothing analogous is available for a generic interior chord, which is exactly why the three tested cases were expected, and found, to fail. But an expectation is not the exhaustive proof a fully general negative claim would need, and this section does not attempt one: it closes the question for the median family (positively) and for every case it actually tested (negatively), and leaves the remainder exactly that narrow, exactly that honestly.

## 26.9 Five more cells, populated from outside this volume's own case studies

Table 26.1's eight populated cells all come from constructions this volume examined at case-study length. That was never the only way a cell could be honestly populated: definition 26.1's four axes classify *any* named construction, not only the ones chosen for a dedicated chapter, and a cell can be filled by a real, citable, correctly-classified projection or grid found in the literature just as validly as by one this volume built a chapter around. This section does exactly that for as many of the sixteen remaining cells as a genuine, sourced search actually turned up a fit for — five, not sixteen — and is explicit, for the rest, about what was tried and why it did not close the cell, in place of either silence or a forced, badly-fitting example.

| Construction | Family | Deployment | Domain | Continuity |
|---|---|---|---|---|
| Conformal cubed sphere (Rančić, Purser & Mesinger, 1996) | Conformal | Global, single instance | Discrete | Contiguous |
| HEALPix (Górski et al., 2005) | Equal-area | Global, single instance | Discrete | Interrupted |
| Albers equal-area conic (1805) | Equal-area | Regional, re-centred | Continuous | Contiguous |
| EASE-Grid, polar sub-grids (NSIDC, 1990s) | Equal-area | Regional, re-centred | Discrete | Contiguous |
| Chamberlin trimetric (1946) | Aphylactic | Regional, re-centred | Continuous | Contiguous |

Table 26.4: Five cells of table 26.1's grid, populated by named constructions from outside this volume's own case-study chapters. None of these five rows is added to table 26.1 itself, to keep that table's own caption — which counts specifically the constructions this volume examined at case-study length — accurate; they are recorded here as a second, explicitly distinct table instead.

**Key Result**

*[Five further cells, each with a real, sourced witness] Table 26.4 places five more named, real, sourced constructions — none of them a subject of a dedicated case-study chapter anywhere in this volume or its companions — into five of table 26.1's previously-unpopulated cells, bringing the count of populated cells from eight to thirteen of the twenty-four definition 26.1's four axes generate.*

*Why each entry classifies the way table 26.4 states. Conformal cubed sphere.* Rančić, Purser

& Mesinger's 1996 construction[4] maps the sphere to the six faces of a circumscribed cube by a mapping constructed, by Schwarz–Christoffel-type methods closely related to chapter 22's own, to be angle-preserving everywhere except at the eight cube-corner singularities — conformal, in definition 26.1's sense, away from a measure-zero set, exactly as Peirce's own quincuncial projection is conformal away from its four singular points. It is deployed as one single global grid covering the whole sphere at once (global, single instance, not a per-face re-centred formula in definition 26.1's sense, since the six faces are six patches of one construction rather than six independent re-applications of a formula to six different centres chosen after the fact); it is discrete by design, built for finite-volume atmospheric and oceanic models rather than as a continuous chart; and it is contiguous as a grid, its six faces sharing edges exactly as chapter 18's ISEA3H grid's twenty faces do, notwithstanding that either grid can be *drawn* as an interrupted net for visualization, a presentation choice section 26.2's continuity axis does not track (the axis classifies the deployed grid, not any one drawing of it, exactly as chapter 18's own ISEA3H row is not reclassified as interrupted merely because its icosahedral net can be drawn flattened and cut).

*HEALPix.* Górski et al.'s Hierarchical Equal Area isoLatitude Pixelization[5] combines, by its own published construction, a Lambert cylindrical equal-area base in the equatorial belt with a distinct, explicitly interrupted Collignon-type construction in each of the four polar diamond regions — equal-area throughout by design (the isolatitude property this scheme is named for exists specifically to make every pixel's area identical), one single global grid rather than a set of separately re-centred instances, discrete, and interrupted by its own polar component's own published description, not by any presentation choice imposed on it here. This is a genuinely different case from the cubed sphere's: where the cubed sphere's apparent interruption is only a drawing convention over a contiguous grid, HEALPix's polar interruption is a structural feature of the projection itself, present in the deployed grid and not only in any one drawing of it.

*Albers equal-area conic.* Heinrich Christian Albers's 1805 projection[6] is equal-area exactly, by the standard construction (a conic projection with two standard parallels chosen so that the compression along meridians exactly cancels the stretching along parallels at every latitude, the textbook argument chapter 5's own trade-off framework already covers as an instance of the equal-area family). It is deployed regionally and re-centred in exactly definition 26.1's sense: the same formula, with its two standard parallels and central meridian reselected for each region mapped, is what underlies its use for the coterminous United States by the USGS and by the U.S. Census Bureau, and equally its reuse, with different standard parallels chosen for the region in question, for continental Europe or for other mid-latitude countries and regions — the same relationship to a single global formula that definition 26.1 already describes for Gauss–Krüger/UTM, with a conic in place of a transverse-cylindrical formula. It is continuous and contiguous, a single connected image of each region mapped,

[4]M. Rančić, R. J. Purser & F. Mesinger, "A global shallow-water model using an expanded spherical cube: gnomonic versus conformal coordinates," *Quarterly Journal of the Royal Meteorological Society* 122(532), 1996, pp. 959–982; cited here, as with several secondary constructions elsewhere in this volume, at the level of its published existence, its stated conformality, and its stated purpose, not from a reading of its body text in this pass.

[5]K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke & M. Bartelmann, "HEALPix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere," *The Astrophysical Journal* 622(2), 2005, pp. 759–771; cited, like the preceding entry, for its published construction and stated properties rather than from a body-text reading in this pass.

[6]H. C. Albers, "Beschreibung einer neuen Kegelprojektion von Europa," *Zach's Monatliche Correspondenz* 12, 1805, pp. 450–459; the projection's equal-area property and its standard modern deployment (two standard parallels reselected per region) are standard cartographic textbook material, given in this form by, among others, Snyder's own USGS reference work, *Map Projections: A Working Manual* (1987), pp. 98–103, not independently re-derived here.

exactly as Lambert's azimuthal equal-area projection (table 26.1's existing equal-area/global row) is.

*EASE-Grid, polar sub-grids.* The Equal-Area Scalable Earth Grid,[7] developed at the National Snow and Ice Data Center for satellite passive-microwave data, is defined on three distinct equal-area projections rather than one: a cylindrical equal-area projection for global and temperate-zone coverage, and a Lambert azimuthal equal-area projection independently centred at each pole for the two polar sub-grids. Table 26.4 records only the two polar sub-grids, not the cylindrical global component, precisely because only the polar pair is genuinely regional and re-centred in definition 26.1's sense — the same Lambert azimuthal equal-area formula, reapplied with its centre shifted from the north pole to the south pole for the second instance, exactly analogous to Gauss–Krüger's zones, rather than a single global instance the way the cylindrical component is. This is worth being as explicit about as section 26.2's own footnote on the ISEA3H DGGS was about a related classification difficulty: the EASE-Grid system *as a whole* does not classify cleanly onto any single row of definition 26.1's grid, because its three component projections do not all share one deployment classification, and splitting it as done here (recording the regional, re-centred polar pair, and leaving the global cylindrical component unrecorded rather than forcing it into either the same row or a further constructed one) is a genuine choice, flagged rather than silently made, not a fact discovered ready-made in the source material. The polar sub-grids are discrete, built specifically for gridded raster data rather than as a continuous chart, and contiguous, each one a single connected raster covering its hemisphere without any interior gap.

*Chamberlin trimetric.* Wellman Chamberlin's 1946 construction for the National Geographic Society[8] is not conformal, not equal-area, and not even equidistant in the strict sense (it minimizes, rather than eliminates, distance distortion from three fixed base points, unlike chapter 8's own two-point-exact construction) — aphylactic by definition 26.1's definition, in the same sense Robinson's compromise projection (table 26.1's existing aphylactic/global row) is. Unlike Robinson, though, it is deployed regionally and re-centred: National Geographic Society atlas maps of individual continents each use their own set of three base points, selected so that the spherical triangle they define encloses most of the continent being mapped, the same triangulated formula reapplied with a different, region-specific choice of base points for Africa, for North America, for Asia, and for every other continent it has been used for — three re-centred points playing the role Gauss–Krüger/UTM's single re-centred meridian plays, rather than the fixed global formula Robinson's own single-instance deployment uses. It is continuous and contiguous, a single connected image of whichever continent's own three base points define its triangle. □

[7] National Snow and Ice Data Center, "A Guide to EASE Grids," nsidc.org, consulted 5 Aug 2026, describing the original (M. J. Brodzik & K. W. Knowles) and version-2.0 (M. J. Brodzik, B. Billingsley, T. Haran, B. Raup & M. H. Savoie, *ISPRS International Journal of Geo-Information* 1(1), 2012, pp. 32–45) definitions; cited at the level of the published grid definition, not independently re-derived here.

[8] W. Chamberlin, *The Round Earth on Flat Paper: Map Projections Used by Cartographers*, National Geographic Society, 1947; the regional, re-centred deployment pattern described here (a distinct triple of base points selected per continent) is documented in A. H. J. Christensen, "The Chamberlin Trimetric Projection," *Cartography and Geographic Information Science* 19(2), 1992, pp. 88–100, cited here at the level of its published account of the projection's deployment history, not independently re-derived.

**Epistemic Notice**

Eleven of the sixteen cells section 26.5 left unpopulated remain unpopulated after this section's search, and it is worth being precise about what kind of negative result that is. This was not an exhaustive search of the cartographic and DGGS literature — a search of that scope is exactly the sort of open-ended undertaking section 26.10 already declines to treat as a closeable todo item elsewhere in this chapter — but a targeted one, aimed specifically at the eleven remaining combinations, including several rounds aimed at the ones that seemed most likely to have a real witness (a regionally re-centred discrete grid, tried against the Rančić–Purser–Mesinger conformal cubed sphere's own regional-refinement literature and against stretched-grid and nested-grid numerical weather models; a discrete aphylactic grid, tried against the DGGS survey literature directly). Two structural observations, not proofs, are worth recording about why several of these combinations may be genuinely rare rather than merely unsearched-for. First, every discrete-and-interrupted combination this section did not fill (conformal/global/discrete/interrupted; every regional/discrete/interrupted combination; the aphylactic/discrete combinations) sits in some tension with discrete global grid systems' own usual motivating purpose, which — as chapter 18's own treatment of the ISEA3H and its companions makes clear — is specifically to provide seamless, non-interrupted global coverage for exactly the kind of statistical and computational applications that a deliberately interrupted grid would complicate; a genuinely interrupted DGGS is not impossible on this evidence, only working against the motivation that has driven most such systems' design. Second, every aphylactic/regional/discrete cell and the aphylactic/global/discrete pair remain unfilled together, and a plausible — again structural, not proved — reason is that discrete grid design has historically optimized directly for one of the two provable extremes corollary 4.5 names (equal-area, for the statistical and remote-sensing applications chapter 18 surveys; conformal, for the numerical-modeling applications section 26.9's own cubed-sphere entry serves), leaving the deliberately-compromise aphylactic family — whose whole rationale is balancing distortions a single-purpose application does not usually need balanced — comparatively unmotivated as a discrete-grid design target specifically, even though named aphylactic *continuous* projections are exactly as real and as old as table 26.1's own Robinson and Fuller rows and table 26.4's own Chamberlin trimetric entry attest. Neither observation is a proof that any of the eleven cells is empty of real constructions; both are offered, in the same spirit section 26.2's own closing disclaimerbox already applies to the original sixteen, as a reason the emptiness found here is not evidence of an oversight, without claiming it is evidence of a genuine structural gap either.

## 26.10 Status of the volume at the close of this draft

With this chapter, all four parts anticipated in section 1.2 and in the preface's running status note contain at least a full first pass: Part I's foundations, Part II's four impossibility theorems (one, section 5.6, still resting on secondary-sourced attribution rather than a primary-verified proof of Milnor's optimality argument specifically), Part III's now eight historical case studies across seven chapters plus one non-historical construction (the compromise-family gap of section 26.2 closed by chapter 19, the interruption gap of section 26.3 closed by chapter 20 and its Mollweide half completed in full closed form after this same drafting pass, the aphylactic/interrupted cell those two closures left open closed in turn by chapter 21, and the remaining conformal/interrupted cell examined but not closed by chapter 22's historical case, then closed by chapter 23's own constructed witness), Part IV's new bridge chapter formally connecting proposition 18.1 to chapter 6 (chapter 25), and this closing taxonomy. What remains, catalogued rather than hidden across the todoboxes of chapters 5, 14, 16, 18 to 23 and 25 and this chapter, is: direct consultation of Milnor's 1969 primary text (its precise formalization is now cross-confirmed by a second, independent secondary source, appendix A.2, but the paper itself remains unread; section 5.8 has since closed the uniqueness-and-regularity portion of this item for Milnor's Chebyshev exposition specifically, via a detailed secondary restatement of Milnor's own hypotheses, section 5.8, leaving Milnor's

own geodesic-disk theorem, theorem 5.10 itself, exactly as unread as before); direct consultation of Mercator's 1569 map legend (Legend 3's own statement of method has since been read in a 1932 English translation reached at one remove, section 14.3, confirming rather than merely reporting that Mercator's text states no derivation, though the actual procedure behind his table of spacings remains as open as before and the Latin original and the remaining fourteen legends are still unread), Lambert's 1772 memoir (since read in full via the 1894 Wangerin critical edition, narrowing but not confirming the secondary-source claim that Lambert explicitly stated the conformal/equal-area exclusivity, section 14.5), Krüger's 1912 monograph (an open-access institutional digitization has since been located and reached, DOI `10.2312/GFZ.b103-krueger28`, GFZ Potsdam; automated text extraction returned only page headers, not the body text, so the primary derivation itself remains unread, section 16.3), Robinson's 1974 publication, Goode's 1925 publication (whose 1925 date places it in the American public domain as of 2021, a concrete access lead recorded in appendix A.7 but not yet followed), Fuller and Sadao's 1954 documentation, still read only through secondary sources, and Gray's 1994/1995 transformation-equation papers, whose own typeset text likewise remains unread but whose specific 1995 construction has since been reconstructed from Gray's own posted source code rather than left to secondary paraphrase alone, as detailed below (Peirce's own 1879 paper, previously on this list with only an access path recorded in appendix A.5, has since been retrieved and read in full, per section 22.11 below); a full reading of the Sahr–White–Kimerling (2003) survey's body rather than its abstract — closed outright by appendix A.4's later reading of all thirteen body pages via a second open mirror (appendix A.4), which also produced, as an unplanned by-product, appendix A.4's partial corroboration of chapter 21's own Gray-reconstruction numbers; that corroboration is itself only doubly-secondary, tracing to White, Kimerling, Sahr & Song (1998) rather than to an independent computation, so a new, narrower item replaces the closed one: White et al. (1998) itself remains a confirmed, explicit publisher paywall rather than a technical access failure, though appendix A.6 has since mined its own abstract for methodological detail (five compared transformation methods, not the four chapter 18 had occasion to name, and exact cell counts up to 65 536 per face) no citing source in this volume had previously recorded (appendix A.6); independent verification of chapter 19's computed distortion figures against Ipbüker's (2004) published numerical evaluation of the same projection; a direct computation of $(a, b)$ for Fuller's own per-face transformation, replacing chapter 21's honestly-flagged gnomonic proxy — no longer outstanding in the form just stated: section 21.5 has since reconstructed Gray's exact construction from his own posted source code and coordinate table (not the typeset paper) and computed its local distortion directly (proposition 21.6), and appendix A.4's independent reading of Sahr, White, and Kimerling's full body text corroborates the reconstruction's qualitative distortion ranking against gnomonic projection (appendix A.4); what remains open is narrower and purely bibliographic — Gray's 1994 and 1995 papers themselves, as typeset journal text, still unconsulted; a primary source for the cited index-holonomy correspondence chapter 25 relies on without re-deriving; a primary source for the general Schwarz–Christoffel construction itself (Christoffel 1867, Schwarz 1869) — narrowed to Christoffel's 1867 paper alone, Schwarz's 1869 paper now read in full (section 22.12), confirming rather than correcting section 22.9's exponent law and section 22.3's lemniscatic-integral formula, and identifying, as an unplanned by-product, that Schwarz's own $\Psi(u, t)$ is the first appearance of what is now called the Schwarzian derivative — though section 22.11 has since established that this citation was never Peirce's own in the first place: his 1879 paper, now read in full, names no source for the disk-to-square step, so the open item is a primary source for the general theory a modern reader would invoke to classify Peirce's map, not for anything Peirce himself cited (section 22.11); section 22.12 had located both papers directly, with two different access outcomes: Christoffel's 1867 paper is open-

access on Zenodo but returns as an unreadable scanned image (section 22.12), still unread, while Schwarz's 1869 paper, reached via the same Internet Archive scan and since supplied directly as a page-image scan, is now closed (section 22.12); Peirce's own 1879 paper, by contrast, is no longer on this list at all, having been retrieved and read in full by that same later pass, which independently confirmed the elliptic-function formula this volume already had at second hand (section 22.11), settled the 394–396-versus-394–397 page-range discrepancy in favor of 394–396 (section 22.11), and discovered that the previously recorded 9 %/13 %/50 % scale-exaggeration comparison attributed to Peirce is not in this text at all, redirecting that specific open question toward Peirce's separate and still-unlocated 1877 Coast Survey report; the branch exponent at Peirce's four singular points, $\alpha = 1/2$, was already computed by section 22.9 by an argument this volume shows does not depend on which exact historical formula is correct, and stands unaffected by this section's closure; a rigorous asymptotic analysis, as $n \to \infty$ and under either convention section 25.7.2 tested, of the pentagon-corner defect's empirically observed decay and of the observed spreading of curvature onto the hexagon-only vertices (section 25.7.2) — narrowed further, though not closed, by section 25.7.3's extension to nine data points spanning $n = 2$ to 20, whose successive log-log slopes approach $-2$ monotonically without a fitted power law or derived rate (section 25.7.3), and narrowed once more by section 25.7.4's actual least-squares fit extending the data to $n = 40$: a restricted large-$n$ fit gives slope $-2.079 \pm 0.007$, a 95% confidence interval that does not quite contain the conjectured $-2$, leaving open whether the true limit is exactly $-2$ approached slowly or a nearby but distinct exponent (proposition 25.12 and section 25.7.4), narrowed once more by section 25.7.5's extension to $n = 50$: the restricted-range slope moves to $-2.049 \pm 0.005$ over $n = 20$ to 50, monotonically closer to $-2$ at every successive restriction tested, still without its confidence interval reaching $-2$ exactly (proposition 25.13 and section 25.7.5); narrowed far more sharply by a genuine derivation rather than a further fit: section 25.7.7 proves outright, from the exact vertex count $V(n) = 20n^2$ (proposition 25.16) combined with the total-defect identity already exact at every finite $n$, that the *mean* defect over all of $\mathcal{P}_n$'s vertices decays as exactly $36^\circ/n^2$ — not fitted, not extrapolated — which, combined with the exact-denominator ratio $\kappa(n) := \delta_{\text{pent}}(n)/\bar{\delta}(n)$ this proof makes possible (definition 25.17), reduces the pentagon-corner exponent question to whether $\kappa(n)$ converges to a finite nonzero limit as $n \to \infty$ (section 25.7.7); computed and fit at fourteen frequencies from $n = 2$ to 50 and then extended, across three further rounds of computation, to $n = 60, 70, 80, 100$ and finally to $n = 150, 200$, $\kappa(n)$'s quadratic- and cubic-in-$1/n$ extrapolations converge, independently of each other and increasingly independently of range, to $\kappa_\infty \approx 0.60282$, stable to five significant figures across every fit tried with $n \geq 30$ reaching at least $n = 100$ (section 25.7.28) — evidence for the exponent being exactly $-2$ now resting on an exact law for the mean plus a tightly bounded, repeatedly re-derived ratio, rather than on a single fitted log-log slope whose confidence interval had not yet reached $-2$ after any of propositions 25.12 and 25.13's successive restrictions; what remains open is narrower than before but not closed: no proof that $\lim_{n\to\infty} \kappa(n)$ exists, and no closed-form identification of $\kappa_\infty$'s value, eighteen icosahedrally-motivated candidates and every rational with denominator under 100 having been tried and rejected (propositions 25.37 and 25.43), with a genuine derivation, per section 25.7.20's todobox, requiring a continuum-limit analysis of the relative curvature-density profile $h(v)$ near an icosahedral vertex — precisely the same profile the next item's shape-statistics findings independently motivate, so that what were two separate open questions in earlier drafts of this list (the pentagon-corner decay rate, and the limiting shape of the defect distribution) are now visibly the same open question asked from two directions, not yet answered from either; narrowed once more, and split into two more specific pieces, by section 25.7.24's independently computed vertex-area ratio $\kappa_{\text{area}}(n)$, which requires no perturbative curvature literature at all and extrapolates cleanly

to $\kappa_{\text{area},\infty} \approx 0.57976$ (section 25.7.24): a pure leading-order Gauss–Bonnet argument (defect proportional to a vertex's dual-cell area alone) now accounts for roughly 96 % of $\kappa_\infty$'s departure from 1, without conjecture, while a genuine and reproducible residual gap of about 0.023 (some 4 % of $\kappa_\infty$) remains unexplained by area alone, redirecting the closed-form question from an undifferentiated "why 0.603" to the narrower "what shape- or anisotropy-dependent correction, beyond dual-cell area, accounts for the remaining 0.023"; narrowed once more, and in an unexpected direction, by section 25.7.25's direct test of the most natural such correction, cell eccentricity: a real correlation between local shape and local defect-per-area does exist among hexagon-only cells (section 25.7.25), but the sixty pentagon-corner cells carry zero shape variance at any tested frequency (machine-precision icosahedral symmetry, section 25.7.25), so no shape-based correction has anything to act on at the pentagon side of the ratio that defines $\kappa(n)$, and fitting the correction in fact closes under 0.2 % of the residual (section 25.7.25); the open item is accordingly restated once more, from a continuous shape-or-anisotropy correction to a specifically *discrete, fivefold-valence-dependent* one — closer in kind to the disclination literature section 25.7.21 already located than to an ordinary per-cell shape statistic; narrowed once more, and ruled out in a different direction, by section 25.7.26's direct test of a third candidate — that section 25.7.24's 0.023 residual might be, at least in part, an artifact of measuring vertex area on a best-fit *plane* rather than on the sphere itself, given section 25.7.1's independently established hexagon non-planarity: recomputing the same ratio with each face's exact spherical (Girard) area in place of the flat SVD-projected one leaves the extrapolated limit unchanged to five significant figures (0.579735 against 0.57976, section 25.7.26), a clean negative result rather than an inconclusive one, since proposition 25.41 independently confirms the spherical measurement itself is exact at every $n$; two of the most natural non-perturbative candidates for the residual — a shape correction, and a curved-versus-flat measurement artifact — are therefore now both ruled out directly, sharpening rather than merely repeating the conclusion that whatever remains is tied to fivefold combinatorial valence under chord-based discretization specifically, addressable, on present evidence, only by the still-unobtained Feinberg–Friedberg–Lee–Ren machinery; and an explicit identification of a specific smooth vector field on the round sphere whose zero set the discrete singularities of chapter 25's combed-flat field converge to as $n \to \infty$, a narrower and now more precisely stated question than the general limiting construction flagged in earlier drafts of this item, since section 25.7 has separately shown that the surfaces $\mathcal{P}_n$ converge to $S^2$ in Hausdorff distance and that the total-curvature identity needs no limit at all for a genuinely flat-faced polyhedron (it is exact at every finite $n$), isolating the singularity-convergence question as the one part still genuinely open — and section 25.7.6 has since given that isolated question a sharper, and rather different, shape of its own: computing three shape statistics of the full (not just pentagon-corner) defect distribution out to $n = 50$ finds each approaching a plateau well short of the values exact uniformity would require, evidence (not a proof) that the right limiting object is a fixed, non-constant relative curvature-density profile on $S^2$ rather than a finite zero set at all, which would mean the original question, as posed, is not merely open but aimed at the wrong kind of limit (section 25.7.6). Eight items from earlier drafts of this list are removed rather than carried forward: the interrupted-aphylactic gap (closed by chapter 21), the Mollweide citation gap (closed within chapter 20 itself), the unexamined status of the conformal-interrupted cell (examined by chapter 22), whether a genuinely interrupted conformal projection can be constructed at all (answered, in the specific narrow sense section 23.4 states carefully, by chapter 23), the uncomputed branch exponent at Peirce's four singular points (computed in section 22.9), the general, unstated form of the discrete-to-continuous limiting question (replaced by section 25.7's sharper items, several closed and one narrowed but still open), the uncomputed sixty pentagon-corner defect values themselves (computed, for $\mathcal{P}_2$ under

one stated convention, in section 25.7.1, which in the course of computing them also discovered and measured the hexagon non-planarity responsible for the convention-dependence now recorded as a new, more specific item above), and the centroid-fan convention's implicit face-centroid defect, left uncomputed in an earlier draft of this list and closed exactly in section 25.7.49 (proposition 25.63 and section 25.7.49). Two further items are narrowed rather than removed by this same drafting pass: the convention-dependence question is substantially, though not completely, resolved by section 25.7.2's two-convention agreement, and the higher-frequency question is answered with three concrete data points by table 25.1 rather than left as an unattempted extension — both are carried forward above only in the sharper, partly quantified form section 25.7.2 leaves them in, not in their original, more open-ended form. None of the remaining open items contradicts anything proved in this draft; each is a specific, named extension of it, in keeping with the working method chapter 1 committed to at the outset.

A further pass (2 Aug 2026), asked to examine chapters 22 and 25's two remaining large blocks of open items directly, closes one genuine residual and re-confirms, rather than newly resolves, several source-access blockers.

*Closed.* The residual, secondary item this list has carried since section 25.4 first closed the primary citation gap for the index-holonomy correspondence — a second, structurally independent proof of $\mathrm{index}(V, v) = \delta(v)/(2\pi)$, distinct from proposition 25.1's unfold-and-transport argument and from Crane, Desbrun & Schröder's own (convergent) derivation — is now supplied: proposition 25.2 reproves the full correspondence via the global Gauss–Bonnet theorem applied to a punctured disk around $v$, combined with an elementary, local computation of a cone circle's geodesic curvature ($\kappa_g = 1/r$, length $\Theta r$, giving $\oint \kappa_g \, ds = \Theta$ directly, without integrating the parallel-transport ODE), plus the standard Riemannian definition of index via winding relative to a parallel frame rather than the discrete-connection formula section 25.4 cited. This is recorded honestly as resting on that one further standard fact (section 25.4), not as a proof from bare first principles, but it satisfies the specific gap the residual todobox named.

*Re-attempted, not closed.* Two access items on this list were re-attempted directly rather than left standing on their earlier record. Schwarz's 1869 paper: a further full-volume transcript fetch, at the largest length allowance available, reached exactly the same stopping point as every prior attempt — ten pages into Christoffel's companion paper, pp. 46–55 of the scanned 1869 volume — confirming section 22.12's own diagnosis that the limitation is a fixed cutoff intrinsic to fetching this resource by this method, not a matter of requesting more length; no free-text route to Schwarz's own p. 105 was found this pass either. Feinberg, Friedberg, Lee & Ren (1984), the paper section 25.7.21 identified as the specific missing machinery for a closed-form derivation of $\kappa_\infty$: a further search for a free copy (author repositories, INSPIRE, arXiv, citing reviews) again found only the paywalled *Nuclear Physics B* abstract page and citations to it, no accessible full text. Both items stand exactly where sections 22.12 and 25.7.21 to 25.7.23 already left them; this pass changes their status from "last attempted several passes ago" to "re-attempted this pass, same result," which is worth recording rather than leaving silently unconfirmed.

*Reviewed, not altered.* Chapter 25's remaining fifteen todoboxes (the $\kappa(n)$ convergence-and-closed-form cascade, the shape-statistics plateau, and the two access items just re-confirmed) and chapter 22's remaining nine were read in full against the chapters' own later content. Every one of them is already exactly as current as it can be made without new primary-source access or a genuine analytic derivation: each earlier, superseded-looking todobox in the $\kappa(n)$ cascade (the $n$=20, $n$=40, and $n$=50 confidence-interval items in particular) is already followed, in place, by the later section that narrows it further, and the chapter's own closing disclaimerboxes (most recently section 25.7.28's) already state the up-to-date position rather than an outdated one. No item was found that could be closed by reasoning alone from material already in this volume; what remains open in both chapters is open for the reasons already stated in place — an unread primary source, or a continuum-limit derivation genuinely not yet carried out — not for lack of a further look.

*Narrowed, in a new direction, by a negative result.* Section 25.7.8 (new in this pass) tests whether $\kappa_\infty \approx 0.60282$ is an invariant of the icosahedral five-fold disclination itself or an artifact of chapter 25's specific single-shot-barycentric-plus-radial-projection mesh ("SSB"), by implementing two further, independently Gauss–Bonnet-validated class-I subdivision schemes ("RMS", recursive midpoint bisection with per-level renormalization; "GEO", recursive two-step slerp) and repeating the same extrapolation. The two alternative schemes converge cleanly to $\kappa_\infty^{\mathrm{RMS}} \approx 0.9237$ and $\kappa_\infty^{\mathrm{GEO}} \approx 0.9520$ — neither matching SSB's 0.60282, nor each other, to better than a few percent (section 25.7.8). This closes off, with a concrete computed counterexample rather than only an a-priori argument, the possibility that a scheme-independent closed form for $\kappa_\infty$ might exist and simply not yet have been found: any future closed-form derivation is necessarily specific to chapter 25's own SSB construction, not to the disclination geometrically. It does not close the closed-form question itself, and opens a new, narrower one in its place — why the two genuinely-geodesic schemes (RMS, GEO) agree with each other far better than either agrees with SSB — which is recorded as a fresh open item in section 25.7.8's own closing disclaimerbox rather than repeated here.

*Narrowed further, same pass: the mechanism behind the gap identified, not yet explained.* Section 25.7.9 (new in this pass) takes up the open item section 25.7.8's disclaimerbox just left — whether SSB's gap from RMS/GEO is driven by its flat-chord interpolation or by its single-shot (rather than nested) normalization — by constructing a fourth scheme, HYB, that isolates the two: GEO's exact two-step recursive construction, but with each `slerp` replaced by a flat linear interpolation plus renormalization. HYB extrapolates cleanly (Gauss–Bonnet validated as before) to $\kappa_\infty^{\mathrm{HYB}} \approx 0.626743$, within ∼4% of SSB and ∼47–52% below RMS/GEO (section 25.7.9)

*Narrowed once more, this pass, in a negative direction: $M(n)$'s bare-fit status is now explained, not merely reconfirmed.* Section 25.7.37 (new in this pass) tests whether $\kappa_\infty^{\mathrm{SSB}}$'s own $O(1)$-per-$n$ localization trick (section 25.7.29, which pushed that constant to $n = 200{,}000$) could do the same for $M(n)$, by tracking the ring-distance from the nearest icosahedral corner of the vertex actually attaining $\max_v \delta(v)$, at $n = 20$ to 300. That distance is not fixed: it grows linearly in $n$, with ratio $r(n)/n$ converging to $1/3$ by $n = 100$ (proposition 25.53) — the maximizing locus migrates toward each face's own centroid rather than sitting near a corner. This rules out, rather than achieves, the localization this item had left as the natural next step, but explains for the first time why $M(n)$ has resisted the treatment that worked for $\kappa_\infty^{\mathrm{SSB}}$: the two are not the same kind of quantity, one a local disclination property, the other governed by flat-chord interpolation error across a face's interior. $\mathrm{CV}(n)$'s own convergence is untouched by this result one way or the other. What remains open is exactly as before, now sharpened for $M(n)$ specifically: no analytic model of interior flat-chord interpolation error has been attempted, and would be needed for any future closed-form or convergence-rate argument that does not reduce to a numerical fit.

*Closed, conditionally: the analytic model just called for is now supplied, and yields a genuine closed-form candidate.* Section 25.7.38 (new in this pass) computes exactly the interior flat-chord interpolation error section 25.7.37's todobox asked for — the SSB scheme's area-distortion Jacobian $J(\alpha, \beta, \gamma) = |h|/r^3$ under central projection, proved in closed form with no appeal to any large-$n$ limit (proposition 25.54). A symmetry argument identifies its maximum exactly at the face centroid, where $\hat{n} \parallel (A + B + C)$ collapses the computation to $M_\infty^{(\mathrm{area})} = 15(7\sqrt{3} - 3\sqrt{15})/(2\pi) \approx 1.2065670269$ (section 25.7.38). This is conditional on the same leading-order area hypothesis already used, with only partial success, for $\kappa_\infty^{\mathrm{SSB}}$ — but here, unlike there, the numerical residual shrinks with $n$ rather than settling at a stable few percent (section 25.7.38), and an independent direct fit of $M(n)$ itself, extended to $n = 400$, lands within its own noise floor of the closed-form value (proposition 25.55). No rigorous bound on the residual is supplied, and the disclaimerbox following section 25.7.38 is explicit that the same clean-looking numerical agreement misled the corresponding $\kappa_\infty$ hypothesis before a later, more careful test found its residual anyway — so this is recorded as a strongly evidenced candidate closed form for $M_\infty$, not a proof, with a concretely scoped next step (a second, centroid-anchored localization, distinct from section 25.7.29's corner-anchored one) named for closing the gap. $\mathrm{CV}(n)$ has no analogous candidate: it is a whole-distribution statistic with no single localized argmax for this kind of argument to attach to, and its convergence remains exactly as open as $\kappa_\infty^{\mathrm{SSB}}$'s own closed form.

# Appendix A

# A Second Verification Pass

## A.1 Purpose of this appendix

Section 26.10 catalogues, rather than hides, every point in this volume where a claim rests on a source consulted only secondhand or not at all. This appendix does not close that list — primary texts remain primary texts, and a citation search is not a substitute for reading Mercator's 1569 legend or Milnor's 1969 proof directly — but it does what the same working method has done at several points earlier in this volume's drafting (most visibly chapter 20's numerical reproduction of the historical join parallel): it takes a second, more careful pass at specific bibliographic and secondary-source detail for a handful of the open items, replacing a general citation with a precise one, and in one case narrowing an open numerical question rather than closing it outright. Each section below is explicit about exactly how much ground was, and was not, covered.

## A.2 Milnor, 1969: precisely identified, and precisely described

Section 5.6's original account of theorem 5.10 rested on two independent secondary summaries of Milnor's optimality argument, without a precise bibliographic identification of Milnor's own paper. That identification is now available: the paper is John Milnor, "A Problem in Cartography," *The American Mathematical Monthly*, volume 76, number 10 (December 1969), pages 1101–1112 — confirmed independently across the journal's own publisher listing, a widely cited bibliography entry in a 2004 *Cartography and Geographic Information Science* survey, and the Mathematical Association of America's own note identifying Milnor as an MIT-affiliated author at the time of publication (consulted 1 Aug 2026; the paper itself remains unconsulted).

**Key Result**

*Beyond the bibliographic identification, several independent secondary sources consulted in this pass converge on a specific description of Milnor's argument that sharpens section 5.6's original, more general summary: Milnor defines a conformal map's distortion as the oscillation (difference between maximum and minimum) of the logarithm of its local scale factor $\sigma$, and shows, in a construction attributed to Chebyshev's 1856 principle of prescribing zero distortion along a closed boundary curve, that the azimuthal equidistant projection minimizes this oscillation for a spherical cap specifically — a Chebyshev-type optimality result for the conformal case, distinct from but structurally parallel to theorem 5.2's own Gauss–Bonnet-derived bound, which this volume proved independently for the equal-area case. This is a more precise statement than section 5.6's original account gave, and is now recorded as such.*

**Key Result**

*[A second, independent source converges on the same formalization] A 2025 preprint by Miyachi, Ohshika, and Papadopoulos on conical map projections (consulted 1 Aug 2026, itself a technical paper by working differential geometers rather than an expository summary) quotes Milnor's own definitions directly rather than paraphrasing them at a distance: for a $C^1$-embedding $f$ of a spherical domain $\Omega$ into the plane and a point $x \in \Omega$, it defines $M(x) = \sup_v |df(v)|/|v|$ and $m(x) = \inf_v |df(v)|/|v|$ over nonzero tangent vectors $v$, the pointwise bi-Lipschitz constant $\sigma(x) = \max\{M(x), m(x)^{-1}\}$, and then $L(f) = \sup_x \sigma(x)$, $\ell(f) = \inf_x \sigma(x)$, with Milnor's "metrical distortion" given by $v(f) = \log(L(f)/\ell(f))$. This is, up to the identification of $\sigma$ with the earlier paraphrase's "local scale factor," the same quantity as the oscillation-of-$\log\sigma$ description above ($\log(L/\ell) = \sup\log\sigma - \inf\log\sigma$), so the two independently consulted secondary accounts corroborate each other on the precise formalization, not merely on the qualitative claim. The same source states, again attributing it to Milnor's 1969 paper without reproducing the proof, that azimuthal equidistant projection achieves the minimum of $v(f)$ among conformal maps of a spherical cap — consistent with, and slightly more specific than, the Chebyshev-type framing above.*

**Todo — verify primary source**

This section narrows, but does not close, section 26.10's Milnor item. What remains unconsulted is Milnor's own derivation — specifically, the proof that the azimuthal equidistant projection achieves the stated optimum, as opposed to the now well-corroborated *statement* that it does. None of the three additional secondary sources consulted in this pass (a SIAM Review bibliography entry, a 2023 expository blog post working through the distortion definition in elementary terms, and the 2004 CaGIS survey bibliography) reproduces Milnor's proof in enough detail for this volume to treat the optimality argument as independently checked; each corroborates the same claim without supplying the derivation. The relationship between *theorem* 5.10's optimality half and *theorem* 5.2's Gauss–Bonnet bound, flagged as open in chapter 5's own closing section, is likewise unaffected by this pass.

**Key Result**

*[Milnor's own text, precisely located — and precisely why this pass still does not read it] A further pass searched directly for the Taylor & Francis paywall's alternatives, rather than for further secondary paraphrase of the same result, and found one specific enough to record.* A Problem in Cartography *was reprinted, without revision beyond Milnor's own added commentary, as the third paper in* John Milnor Collected Papers, Volume I: Geometry *(Publish or Perish, 1994; reissued by the American Mathematical Society), at pages 159–172 — a pagination confirmed directly against Milnor's own bibliography, hosted at his Stony Brook faculty page (`math.stonybrook.edu/~jack/milnor-pub.pdf`, consulted 1 Aug 2026), which lists the entry as `[c] A problem in cartography. Amer. Math. Monthly, 76: 1101--1112; [CP-1, 159--172]`. The volume itself is digitized and hosted by the Internet Archive (identifier `collectedpapers0001miln`, 322 pages, English, OCR already run with reported confidence 1.0), confirmed by direct retrieval of the item's own catalogue page in this pass. That same retrieval, however, returned the item's own `Access-restricted-item: true` flag: the scan sits in the Archive's controlled-digital-lending collection (tagged `inlibrary` and `printdisabled`), requiring an Archive account and an active one-copy loan to read even a single page, not an openly fetchable transcript of the kind sections 16.3 and 22.12 obtained for Krüger's monograph and the 1869 Crelle volume. No borrowing session was initiated in this pass, consistent with this volume's standing practice of not attempting access routes that require credentials or accounts this working method does not hold.*

**Todo — verify primary source**

Appendix A.2 narrows section 26.10's Milnor item in a specific, useful way — a precise edition, page range, and a confirmed digitization are now on record, where before only the original 1969 journal citation and its paywall were known — but it does not close the item, and the reason it does not is now precisely characterized rather than left as a generic "behind a paywall" note: this is a controlled-lending access barrier, not an open-access one, and closing it would require either a library that holds a physical or unrestricted-digital copy of the 1994 *Collected Papers* volume, or an Internet Archive borrowing session this pass did not undertake. Section 5.8's todobox items (a) and (c) — the optimality proof's exact technique, and Milnor's own stated relationship, if any, between $\delta_0(\alpha)$ and total curvature — accordingly remain exactly as open as before; only the map of how one might eventually reach them has improved.

## A.3 Ipbüker, 2004: precisely identified, and precisely described

Section 26.10 listed, among its open items, independent verification of chapter 19's computed distortion figures against Ipbüker's published numerical evaluation, cited there without a precise bibliographic identification. That identification: Cengizhan Ipbüker, "Numerical Evaluation of the Robinson Projection," *Cartography and Geographic Information Science*, volume 31, number 2 (2004), pages 79–88, DOI 10.1559/1523040041649425 — confirmed via the journal's own table of contents listing and independently via a bibliography compiled for NASA's G.Projector map-projection software documentation, both consulted 1 Aug 2026.

**Key Result**

*The paper's own abstract and citing literature (consulted, not the paper's body) describe its content more specifically than section 19.3's original citation could: Ipbüker compares four different interpolation algorithms — including a multiquadric radial-basis-function method — for deriving continuous analytical expressions from Robinson's original look-up table, and, using the resulting fitted functions, determines the specific latitude at which the parallel scale factor is exactly true and the specific latitude at which angular distortion reaches zero. This is a different distortion analysis from section 19.3's own, which used cubic-spline interpolation of the same table (following the standard practice reported by chapter 19's secondary sources) rather than Ipbüker's multiquadric method, and asked a different question (the overall distortion ellipse across latitude, rather than the specific zero-crossing latitudes Ipbüker's paper isolates).*

**Key Result**

*[This gap's own structure, now cross-checked: two distinct latitudes, precisely located, close to but not exactly matching an independently corroborated round figure] A further pass, prompted by this section's own description of what Ipbüker's paper computes, went back to chapter 19 and performed the natural next step short of reading the paper itself: root-finding, on chapter 19's own full 5° table, for the exact latitude where $k = 1$ (parallel scale true) and the exact latitude where $\omega = 0$ (angular distortion vanishes), rather than reading approximate values off a coarse excerpt (proposition 19.1 and section 19.3). The result — 37.75°–37.82° for true scale and 32.85°–32.91° for zero angular distortion, depending on interpolation method, with the two methods agreeing with each other to within 0.07° — reproduces Ipbüker's own paper's* structure *exactly (two distinct latitudes, not one, matching this section's own description above) and lands within a quarter of a degree of the 38° figure independently reported by MathWorks' and Esri's Robinson-projection reference documentation (consulted 1 Aug 2026, neither citing Ipbüker directly). This is the closest this volume has come to a genuine cross-check of chapter 19's own numbers against an independent source, though it remains, honestly, a check against a rounded secondary figure rather than against Ipbüker's own multiquadric-fitted result.*

**Todo — verify primary source**

This is a narrowing, not a closing, of section 26.10's item: this pass has identified precisely what Ipbüker's paper computes and how it differs in method from chapter 19's own cubic-spline approach, and appendix A.3 has now gone one step further, computing the same two exceptional latitudes directly and cross-checking them against an independently corroborated rounded figure — but this pass has still not obtained Ipbüker's specific numerical results (the true-scale and zero-angular-distortion latitudes as *he* computed them, to his own stated precision) to check against chapter 19's figures directly. The paper is behind a publisher paywall not accessed in this pass, and no illegitimate access route was attempted. Chapter 19's own numbers therefore remain independently computed and now independently cross-checked against a corroborated secondary figure, but not cross-checked against this specific published primary source, exactly as section 26.10 originally flagged, now with both the comparison's target and the size of the remaining gap (0.18°–0.25°, about three to four times chapter 19's own two-method interpolation spread) precisely specified.

## A.4 Sahr, White, and Kimerling, 2003: the full body, obtained and read

Section 26.10 listed a full reading of the Sahr–White–Kimerling (2003) survey's body, as opposed to its abstract alone, as an open item underlying chapter 18. An earlier pass in this same appendix obtained roughly the paper's first three pages via an openly accessible mirror. A further pass has now obtained and read the complete paper (all thirteen body pages, *Cartography and Geographic Information Science* 30(2), pp. 121–134, including Table 1's sixteen-system design-choice summary, Figures 3, 11, and 12, the paper's own worked example construction of the ISEA3H DGGS, and its full reference list) via a second openly accessible full-text mirror, consulted 1 Aug 2026.

**Key Result**

*[Several claims chapter 18 previously sourced only from the abstract or from citing literature are now confirmed, or corrected, against the paper's own body text] Three points are worth recording specifically. First, the paper's own conclusion confirms, in the authors' own reasoning rather than a citing paper's summary, exactly the design-choice chain section 18.2 already attributed to it: icosahedron chosen for its lower distortion (smaller faces mean less distortion transforming each face to the sphere, a general point the paper makes early and then applies); an equator-symmetric orientation (poles at edge midpoints, not vertices) chosen over the more common vertices-at-poles arrangement; aperture-3 hexagons chosen as the smallest aligned hexagon hierarchy; and the inverse ISEA (Snyder 1992) projection chosen specifically for its equal-area property, with cell centers taken as the associated points. Second, and directly relevant to chapter 21 elsewhere in this volume: the paper's own body identifies exactly four transformation methods that map straight planar face edges to the correct spherical great-circle-arc edges — the gnomonic projection (available for any polyhedron, but with comparatively large area and shape distortion), Snyder's ISEA projection (equal-area, but with more shape distortion and less regular cell edges than a small-circle-subdivision alternative), an implementation of Fuller's Dymaxion projection specifically attributed to Gray's 1995 paper, and Dutton's ZOT projection for the octahedron — and states explicitly, as the authors' own finding rather than this volume's, that the Fuller/Gray projection has less area and shape distortion than the gnomonic projection and less shape distortion than the ISEA projection, while not itself being equal-area. Third, the paper records Fuller's own Dymaxion orientation in precise, reproducible form: one icosahedron vertex at 5.2454° W longitude, 2.3009° N latitude, with an adjacent vertex at azimuth 7.46658° from the first — a specification this volume had not previously recorded for either chapter 18 or chapter 21.*

**Key Result**

*[An unplanned, independent cross-check of chapter 21's own numerical finding] Appendix A.4's second point bears directly on proposition 21.6's own reconstruction of Gray's exact equations, computed independently and numerically in that chapter without reference to this survey: this volume's own computation found the Fuller/Gray construction to have substantially* less *area distortion than the gnomonic proxy at both sample points tested (*1.053 *vs.* 1.227 *at the edge midpoint;* 1.134 *vs.* 1.993 *at the vertex), a qualitative finding this appendix's reading of Sahr–White–Kimerling's own body text now confirms independently, from a peer-reviewed source that examined the same comparison (Fuller/Gray against gnomonic) for a different purpose (choosing a DGGS transformation method) and reached the same qualitative ranking. This is a genuine, if partial, independent verification of chapter 21's own reconstruction — partial because Sahr, White, and Kimerling's own comparison is qualitative (a ranking, not the specific numerical figures this volume computed), and because their claim traces to White et al. (1998), a paper this volume has not itself consulted, rather than to an independent computation of their own. It nonetheless narrows the concern, flagged in chapter 21's own todobox, that definition 21.4's reconstruction might not faithfully capture Gray's actual construction: an independently reasoned, independently motivated source agrees with this volume's own reconstruction on the direction of the effect.*

**Todo — verify primary source**

This closes the specific item section 26.10 listed — the paper's body, as opposed to its abstract, has now been read in full, not merely narrowed further — but two smaller items remain. First, the mirror consulted is a third-party host, not the journal's own publisher platform, so page-image fidelity (figure quality, any errata in the original typesetting) has not been cross-checked against the publisher's own version. Second, appendix A.4's cross-check traces Sahr, White, and Kimerling's own Fuller/Gray-versus-gnomonic distortion comparison to White et al. (1998), *Comparing area and shape distortion on polyhedral-based recursive partitions of the sphere*, *International Journal of Geographical Information Science* 12: 805–27 — a paper this appendix has not obtained or read, and which would be the natural next primary source to consult for a fully independent, rather than doubly-secondary, verification of chapter 21's own numbers.

## A.5 Peirce (1879) and Schwarz–Christoffel: an access path found, two corrections precisely dated

Section 26.10 listed both Peirce's own 1879 paper and a primary source for the Schwarz–Christoffel construction chapter 22 cites without re-deriving. Neither is closed by this pass, but both are narrowed further than chapter 22's own footnotes had left them. (A later pass did go on to retrieve and read Peirce's 1879 text itself, closing the first of these two items outright; section 22.11 records that pass and its findings, including the discovery that Peirce's own text names no Schwarz–Christoffel source, so this appendix's second item — a primary source for the general Schwarz–Christoffel construction, as opposed to Peirce's own paper — remains open on its own terms rather than being closed by the same retrieval. A still later pass, recorded within chapter 22 itself rather than in this appendix, located and fetched both papers directly: Christoffel's 1867 paper proved genuinely open-access but unreadable under automated extraction, while Schwarz's 1869 paper was narrowed to reachable-and-plausibly-legible, pending only a further fetch of its own page range (section 22.12). Neither fetch has yet been used to check anything this appendix or chapter 22 derived from Driscoll and Trefethen's modern secondary account, so the item remains open in substance even as its access route is now precisely mapped.)

**Key Result**

*The Schwarz–Christoffel attribution is confirmed, with a precise year for each independent discovery, by a source distinct from the ones chapter 22 originally consulted: Christoffel's contribution in 1867 and Schwarz's in 1869, cited consistently across an arXiv preprint on the mathematics of Peirce's projection and Wikipedia's own account of the same construction. Separately, this pass located precise citations for the two historical corrections chapter 22's own footnote already anticipated qualitatively ("later shown by Pierpont …to require a minor correction"): James Pierpont, "Note on C. S. Peirce's Paper on 'A Quincuncial Projection of the Sphere'," American Journal of Mathematics, volume 18, number 2 (1896), pages 145–152; and I. Frischauf, "Bemerkungen zu C. S. Peirce Quincuncial Projection," the same journal, volume 19, number 4 (1897), pages 381–382. Neither had a precise citation in chapter 22 before this pass.*

**Key Result**

*[A free access route located for both the 1879 paper and the 1896 Pierpont note, distinct from the JSTOR paywall and the NOAA scan this volume had already tried] A further pass, prompted specifically by the JSTOR access wall already recorded above, located a third route neither previous pass had tried: both American Journal of Mathematics volume 2 (1879, containing Peirce's paper at pp. 394–396/397) and volume 18 (1896, containing Pierpont's note at pp. 145–152) fall inside JSTOR's own "Early Journal Content" program, which releases pre-1923 journal content without a paywall, and both volumes are hosted as full page scans with OCR-derived full text on the Internet Archive (`archive.org/details/sim_american-journal-of-mathematics_1879_2` for volume 2; the analogous `sim_american-journal-of-mathematics_1896_18` identifier pattern, unconfirmed for volume 18 specifically in this pass, for volume 18), confirmed by direct retrieval of volume 2's item page and its download options (a `_djvu.txt` OCR transcript and a full-page PDF, both consulted 1 Aug 2026) rather than by a bibliography entry alone. This is a genuinely different fact from anything the earlier JSTOR and NOAA passes established: the earlier passes correctly reported that JSTOR's own reading interface and the NOAA Coast Survey scan were each blocked, but neither pass checked whether the underlying pre-1923 text was separately available through JSTOR's own public-domain release channel.*

**Todo — verify primary source**

This narrows, but importantly does not close, section 26.10's item: a legitimate access route to the primary text is now identified and confirmed reachable (the volume 2 item page and its OCR transcript were retrieved in this pass), but this pass did not go on to retrieve and read the specific transcript pages corresponding to Peirce's pp. 394–396/397 within that transcript, nor did it confirm the parallel Internet Archive identifier for American Journal of Mathematics volume 18 well enough to treat Pierpont's note as similarly reachable rather than merely likely reachable by the same program. A direct follow-up attempt, recorded in chapter 22 alongside section 22.2's encyclopedic-secondary-sourced version of Peirce's own claimed distortion figures, tried to retrieve those specific pages from the transcript located here and did not succeed within this pass's tool limits. The bibliographic discrepancy already flagged (394–396 vs. 394–397 for Peirce's own page range) is consequently still unresolved, since resolving it would require reading the article itself rather than a bibliography entry about it. A future pass that retrieves the specific pages, rather than only confirming the route to them, would close this item outright.

Separately from this appendix, chapter 22 itself has since taken a further, independent pass at the same item, in the manner appendix A.7 below records for chapter 21's own Gray narrowing: returning to a source already cited in that chapter (the same 2016 *Revista Integración* paper this section drew its formulas from) and reading it more closely rather than consulting a new source, it recovered a specific worked example of the Peirce/Pierpont discrepancy at $l = \theta = 0$ (section 22.2) and a second, only partially corroborating data point on the page-range question — a second source agreeing with JSTOR's 394–396 against Arisbe's 394–397, without resolving which is correct (section 22.2). A renewed attempt in that same pass to locate a directly fetchable transcript URL for the Internet Archive item

identified above again did not succeed; the access route remains confirmed, the text itself still unread.

## A.6 White, Kimerling, Sahr, and Song, 1998: precisely identified, still paywalled, and its abstract mined for what it actually claims

Appendix A.4's own closing todobox named this paper as the natural next primary source to consult: appendix A.4's cross-check of chapter 21's Gray-reconstruction numbers against Sahr, White, and Kimerling (2003) is itself only a doubly-secondary corroboration, since Sahr, White, and Kimerling's own Fuller/Gray-versus-gnomonic ranking traces to White, Kimerling, Sahr, and Song (1998), *Comparing Area and Shape Distortion on Polyhedral-Based Recursive Partitions of the Sphere*, *International Journal of Geographical Information Science* 12(8), 805–827. A further pass took this up directly.

**Key Result**

*[The paper's own methodology, now known precisely rather than inferred from citing literature] The publisher's own abstract, retrieved directly from Taylor & Francis rather than from a citing paper's paraphrase, specifies the comparison's design in more detail than any secondary source consulted anywhere else in this volume: five transformation methods were compared side by side — the gnomonic projection, Fuller's Dymaxion projection, Snyder's (1992) equal-area polyhedral projection, a direct spherical subdivision method, and a recursive polyhedral projection — applied to recursive partitions of both the octahedron and the icosahedron, at two density ratios (4-fold and 9-fold) carried respectively to the eighth level (*65 536 *cells per polyhedral face) and the fifth level (*59 049 *cells per polyhedral face). This is new, primary-sourced information for this volume: appendix A.4's own appendix A.4 had already recorded, from Sahr–White–Kimerling's 2003 body text, that White et al. (1998) examined four of these five methods in the specific gnomonic/ISEA/Fuller-Gray/Dutton-ZOT grouping relevant to DGGS construction; the 1998 abstract itself shows the original comparison was broader (five methods, not four, and explicitly including a direct spherical subdivision method Sahr–White–Kimerling's own later summary does not mention), and gives the two density ratios and exact cell counts that no source consulted anywhere else in this volume had recorded.*

**Todo — verify primary source**

This narrows, but does not close, the item appendix A.4's todobox named. The publisher's own abstract page (`tandfonline.com/doi/abs/10.1080/136588198241518`) is reachable and gives the precise methodological detail appendix A.6 records, together with a direct link to the article's PDF marked `needAccess=true` — an explicit, unambiguous paywall, not a technical access failure of the kind that stalled the Krüger and Close items elsewhere in this appendix. A search for an open-access mirror — author repository, ResearchGate full text, the U.S. EPA's own document repositories (Denis White's employer at the time, and a plausible source for an agency-hosted preprint given chapter 18's own unrelated EPA-authored sources), and general web search for a freely hosted PDF — returned only further citations, abstracts, and reference-list entries, never the article's own body text or figures. appendix A.4's cross-check therefore remains exactly as partial as appendix A.4 left it: an independent qualitative corroboration from a source (Sahr–White–Kimerling 2003) that itself only summarizes White et al. (1998)'s findings, not a direct comparison against that paper's own reported numbers. What this pass adds is not access but precision: the comparison's actual scope (five methods, not the four chapter 18 had occasion to name) and its exact scale (cell counts up to 65 536 per face) are now on record, and any future pass attempting this item again has a single named access point (the Taylor & Francis DOI above) rather than only the bibliographic citation.

## A.7 What this pass changes, and what it does not

Neither section above closes an item on section 26.10's list. What changes is the precision with which several of that list's items are now stated: a vague "secondary-sourced account" becomes a specific journal, volume, page range, and description of content, in each case corroborated across more than one independently consulted source rather than one, and in the Milnor and Schwarz–Christoffel cases, cross-checked against a second source found independently of the first. This is worth recording as its own small methodological point, consistent with the standard chapter 1 set at the outset: a todobox that names exactly what is missing and exactly what has been checked so far is more useful to a future revision of this volume than one that only names the missing primary text, and this appendix's contribution is entirely of that sharpening kind, not a substantive mathematical or historical result in its own right.

**Epistemic Notice**

This appendix, across its three drafting passes, has now taken up four of section 26.10's items — Milnor (mathematical, narrowed), Ipbüker (cartographic-technical, narrowed further by appendix A.3's own root-finding cross-check), the Sahr–White–Kimerling survey's body (now closed in full, appendix A.4, with an unplanned by-product: appendix A.4's independent corroboration of chapter 21's own Gray-reconstruction numbers), and the Schwarz–Christoffel attribution underlying chapter 22 (two corrections now precisely dated, and, separately, a free Internet Archive access route to Peirce's own 1879 paper located and confirmed reachable, appendix A.5, though the specific pages still unread) — treating each carefully rather than attempting all of them shallowly. Separately, and outside this appendix's own two sections, chapter 21 itself has since narrowed the Gray 1994/1995 item on section 26.10's list, reconstructing the 1995 paper's specific per-face construction from Gray's own posted source code and coordinate table rather than from the typeset paper (section 21.5), a narrowing this appendix's own appendix A.4 now corroborates from an independent direction. The remaining items — Mercator's 1569 legend, Lambert's 1772 memoir, Krüger's 1912 monograph, Robinson's 1974 publication, Goode's 1925 publication, Fuller and Sadao's 1954 documentation, Gray's 1994/1995 papers as typeset journal text, and Peirce's own 1879 paper — were checked for accessibility in this same pass (an access path, not the text itself, is now recorded for Peirce; a precise citation, not yet a text, for Goode, whose 1925 publication date places it in the American public domain as of 2021, a specific and actionable lead this appendix had not previously recorded) but remain otherwise exactly as open as section 26.10 left them. A further pass since located Goode's specific 1925 publication on HathiTrust (volumes 15–16, full view, University of Michigan copy) but could not retrieve its page images or OCR text directly, since HathiTrust's page-turner interface disallows automated access; chapter 20 itself has, separately, narrowed the $0.2''$ join-parallel residual its own earlier draft flagged, finding that a technical secondary source (Esri's ArcGIS Pro documentation) matches this volume's own independently computed value to the arcsecond, against the more commonly repeated figure traced to Snyder 1993 (section 20.4) — the same access-confirmed-but-text-unread pattern already recorded for Peirce, now added for Goode as well. A still later pass found a second, independent access route alongside the HathiTrust copy already recorded: the paper's actual journal of record is not *Geographical Review* but the *Annals of the Association of American Geographers*, volume 15, number 3 (1925), pages 119–125, DOI `10.1080/00045602509356949`, confirmed via the publisher's (Taylor & Francis) own abstract page and independently via a Semantic Scholar record citing the identical DOI and page range (both consulted 1 Aug 2026; recorded in chapter 20's own footnote alongside this appendix). The publisher page itself sits behind a paywall this pass did not cross, so the text remains unread through this route as well, but the precise journal identification corrects what had been, until this pass, an unspecified periodical reference, and gives any future attempt a single unambiguous target rather than the HathiTrust volume's own less specific listing. A separate pass, working on chapter 16 rather than on chapter 20 or chapter 22, found the same access-confirmed-but-unread pattern a third time, and for a reason distinct from a paywall or a blocked interface: Krüger's 1912 monograph has its own institutional open-access digitization, hosted by GFZ Potsdam with a persistent DOI, genuinely reachable and reached in that pass, but returning only page headers rather than body text under this pass's automated extraction — a scanned-document OCR limitation rather than an access restriction, and recorded in full in section 16.3. A repeat fetch of the same GFZ file in a later pass returned the identical result, confirming the limitation is stable rather than an artifact of a single failed attempt.

A still later pass added two further, small narrowings to this same family of items, neither closing anything but each sharpening precisely what stands between this volume and the primary text. First, appendix A.2 in this appendix located a specific, confirmed digitization of Milnor's own 1969 paper — its 1994 *Collected Papers* reprint, pages 159–172, hosted by the Internet Archive under identifier `collectedpapers0001miln` — and found it to sit behind a genuinely different kind of barrier than the Taylor & Francis paywall already on record: controlled digital lending, requiring an account and an active loan rather than a subscription or a one-time fee, which this pass did not attempt to initiate. Second, section 22.12's own further pass, within chapter 22, extended its earlier partial read of the 1869 Crelle volume's OCR transcript substantially further into the same file — through the opening sections of Christoffel's paper rather than stopping at its first page — without yet reaching Schwarz's own paper at p. 105, narrowing the remaining gap to a matter of retrieval length or starting position rather than of access or legibility.

## A.8 Two items added by the five later chapters

Chapters 8 and 10, added after every pass recorded above, each name one primary source neither previously appeared on this volume's open list, and neither has been pursued by any pass to date: Maurer's 1919 paper introducing the two-point equidistant construction (*Annalen der Hydrographie und Maritimen Meteorologie* 47(3–4), 75–8) and Close's independent 1921 note (*The Geographical Journal* 57(6), 446–448), both reported in chapter 8 from Snyder (1993) and current cartographic reference documentation rather than read directly; and the 1907 papers of Poincaré and Koebe establishing the uniformization theorem chapter 10 states as section 10.4, reported there from Farkas–Kra's textbook treatment. Both are recorded here in the same spirit as this appendix's other entries — named precisely enough that a future pass could attempt them — and neither has yet been attempted; this section adds them to the list rather than closing anything.

### A.8.1 A pass on Close (1921): access route confirmed, article text not yet extracted

A later pass took up the Close half of this pair. *The Geographical Journal* volume 57 (January–June 1921) is in the American public domain (published before 1923) and has been digitized in full by the Internet Archive, identifier `geographicaljour5719unse`, with an OCR full-text transcript openly available at no access barrier — no paywall, no controlled-lending queue of the kind that stalled the Milnor and Peirce items above. This pass fetched that transcript and, from its front-matter table of contents (itself part of the primary scanned volume, not a secondary index), confirmed several details beyond what chapter 8 previously had from Snyder (1993) and Wikipedia's citation alone: Close's note is the first of four grouped short pieces under the running head "Four Notes on Map Projections" in No. 6 (June 1921), its own title is precisely "Note on a Doubly-Equidistant Projection" (not "A Doubly Equidistant Projection" as some later secondary citations render it), its byline reads "Colonel Sir Charles Close, K.B.E., F.R.S., Director-General of the Ordnance Survey," and it occupies pp. 446–448 exactly as already cited — immediately followed in the same grouped set by Hinks on the Allied Western Front maps (p. 448), McCaw on a modified rectangular polyconic projection (p. 451), and Hinks again on projecting the sphere onto a circumscribed cube (p. 454).

This narrows, but does not close, appendix A.8's Close item: the access route is now confirmed genuinely open (unlike Milnor's controlled-lending barrier or Robinson/Ipbüker's publisher paywall), and several bibliographic details are now primary-confirmed rather than secondary-reported, but this pass did not extract or read the three pages of Close's own mathematical content — the OCR transcript for the full 588-page volume is a single large text stream, and locating and retrieving specifically pp. 446–448 of it (rather than the volume's opening pages, which is as far as this pass's single retrieval reached) was not accomplished within this pass. A repeat attempt at a substantially larger retrieval size returned the identical opening portion of the transcript rather than reaching further into the volume, indicating the limitation is on the retrieval side of this pass's tooling rather than resolvable by simply asking for more; a future pass would need either a different retrieval mechanism (e.g. the volume's per-page image files, or a search-inside API rather than the flat OCR transcript) to reach pp. 446–448, not merely a larger request against the same one.

### A.8.2 Maurer (1919): a plausible open access route identified, record not located

A still later pass revisited the claim, made when this item was first added, that Maurer's 1919 *Annalen der Hydrographie und maritimen Meteorologie* paper "is not known to have any comparably open digitization." That claim is now weaker than it was: the journal's institutional successor, the Bundesamt für Seeschifffahrt und Hydrographie (BSH), maintains a public Goobi-based digital library (`digitale-bibliothek.bsh.de`) that has digitized runs of the *Annalen* back into the nineteenth century, with a page-addressable full-text API (`/viewer/api/v1/records/{id}/plaintext/`) that this pass confirmed working on several other volumes and years. This pass did not, however, locate the specific catalog record for volume 47 (1919) among the collection's non-sequentially-numbered record identifiers within this pass's search budget, and so has not obtained or read Maurer's paper itself. This is recorded as a distinct, smaller narrowing from the Close one above: a plausible and likely-open access route is now known to exist, where previously none was known at all, but the specific record has not been found, let alone the three-page article within it read.

### A.8.3 A further pass on Maurer (1919): the catalog's own record numbering resists a shortcut, and the platform's public status is itself now less certain than assumed

A still later pass attempted to close appendix A.8.2's remaining gap directly, by treating the BSH catalog's record identifiers for nearby, already-located *Annalen* volumes as a basis for interpolation: volume 37 (1909) is record `142022`, volume 22 (1894) is `141982`, volume 12 (1884) is `141959`, and volume 57 (1937) is `141848` — a set that, taken together, is not monotonic in either direction against publication year, ruling out the simple guess (tried and confirmed to fail) that volume 47's record would sit at a predictable offset from its neighbors. This pass did, however, independently reconfirm that the plaintext API route itself works as described: fetching `/viewer/api/v1/records/141295/tei/en/` directly, for an unrelated volume encountered while searching, returned the article's own running text (an oceanographic paper on the German Southwest African coast, correctly rendered German prose, not a header-only extraction of the kind that stalled section 16.3's attempt on Krüger). What this pass adds that qualifies rather than simply extends appendix A.8.2's optimism is a direct statement, found on the BSH's own current documentation page for the Digital Library, that the library is not represented as fully public: it states explicitly, in the institution's own English-language wording, that the collection is not yet open to general access and that public activation is still in preparation, notwithstanding the individual records and API routes this appendix and its predecessor pass were able to reach directly by URL.

**Todo — verify primary source**

This narrows appendix A.8.2's item in two directions at once, one encouraging and one cautionary, and closes neither. Encouraging: the plaintext API is confirmed working end to end on at least one genuine record, not merely reachable as an endpoint, so if volume 47's specific record number is ever found, extracting Maurer's three pages is very likely mechanical rather than blocked by an OCR or scanning failure of the kind seen elsewhere in this appendix. Cautionary: the institution's own public statement that the library is not yet officially open raises a question this volume had not previously asked of any of its access routes — whether reaching individual records directly by guessed or found URL, as this appendix has now done twice, constitutes the kind of access this volume should rely on and cite as open, given that the host institution itself does not yet represent the collection as public. This volume takes no position on that question beyond recording it; a future pass attempting to close the Maurer item outright should resolve it, together with the still-unresolved task of locating volume 47's specific record identifier, before treating any eventual text as read via a route this volume can straightforwardly stand behind.

### A.8.4 A pass on Poincaré (1907) and Koebe (1907): precise citations, one open discrepancy, access blocked mid-pass

A further pass took up the Poincaré/Koebe half of the same two-item addition, recorded directly in chapter 10's own todobox (section 10.4) rather than only here. It established precise citations for both papers where the volume previously had none beyond the secondary Farkas–Kra textbook reference: Poincaré's *Acta Mathematica* 31, pp. 1–63, and the first of Koebe's three 1907–1908 *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen* communications, pp. 191–210. It also confirmed, from International Press's own current statement about the journal, that Acta Mathematica's full historical archive is openly accessible by arrangement with the Institut Mittag-Leffler — a materially different situation from Milnor's controlled-lending barrier or Ipbüker's and Goode's publisher paywalls. This pass's own attempt to fetch the article through Project Euclid's hosting, however, was stopped by an automated bot-protection challenge rather than reaching any access-restriction page, so the confirmation of openness rests on the publisher's general statement rather than on this pass's own successful retrieval of the specific article. A small bibliographic discrepancy surfaced in the process and was not resolved at the time: some secondary sources give Poincaré's closing page as 64 rather than 63, a gap of exactly the same flavor as appendix A.5's already-recorded Peirce page-range discrepancy (394–396 vs. 394–397). Koebe's own two later 1907–1908 communications (pp. 633–669 and pp. 337–358) are named but not pursued at all by this pass. Neither derivation has been checked against this volume's own statement of theorem 10.3 and section 10.4.

### A.8.5 The 63-versus-64 discrepancy resolved, from a third independent source

A later pass resolved the page-range discrepancy just recorded. The German Wikisource project maintains an independently-compiled, self-described "vollständig erfasst" (completely captured) table of contents for the entire run of *Acta Mathematica*; its entry for Band 31 (1908) lists "H. Poincaré: Sur l'uniformisation des fonctions analytiques, S. 1–63." This is a third source, compiled independently of both Snyder's secondary account and the Farkas–Kra textbook citation already used in chapter 10, and it agrees with them against the lone outlier giving 64. On the ordinary evidentiary standard already applied to appendix A.5's Peirce page-range question — where multiple independent secondary sources agreeing was treated as sufficient short of reading the primary text — this closes the 63-versus-64 question. It does not, however, substitute for reading Poincaré's own text: the article itself remains

unretrieved, still blocked by Project Euclid's Incapsula challenge, so the substance of the derivation is still reported from Farkas–Kra alone.

## A.9 Two items added by chapter 11

Chapter 11, added after every pass recorded above, names two primary sources neither previously on this volume's open list, in the same spirit as appendix A.8's own Maurer/Close/Poincaré–Koebe additions: David Hilbert's 1901 paper stating and proving theorem 11.8 itself (*Transactions of the American Mathematical Society* 2(1), 87–99), and J. N. Hazzidakis's 1880 paper supplying proposition 11.6's area formula (*Journal für die reine und angewandte Mathematik*, the volume and page range for which this pass has not confirmed). Neither has been pursued by any pass to date; both are recorded here precisely enough that a future pass could attempt them, following exactly the access-route methodology appendix A.8.1 used successfully for a comparably aged *Geographical Journal* volume. Hilbert's paper is a particularly promising candidate for that same treatment: a 1901 *Transactions of the AMS* volume falls inside the same pre-1923 American-public-domain window that made Close (1921) and Goode (1925) tractable elsewhere in this appendix, and *Transactions of the AMS* volumes of this vintage are, on the pattern established by every other pre-1923 English-language journal this appendix has checked, a plausible target for an open Internet Archive or JSTOR Early Journal Content digitization — a plausibility this pass records but has not itself confirmed by a search.

**Key Result**

*[Hilbert (1901) closed, in the published original — the promising candidate confirmed] A later pass confirmed the plausibility flagged just above and located the paper directly: David Hilbert,* Ueber Flächen von Constanter Gaussscher Krümmung, Transactions of the American Mathematical Society *2(1) (1901), pages 87–99, is held in the Internet Archive's* `jstor_ejc` *collection (identifier* `jstor-1986308`*), part of JSTOR's Early Journal Content program, which places pre-1923 journal articles in unrestricted open access with no account, login, or lending session required — confirmed directly by fetching the item's own full-text transcript in this pass, not merely its catalogue page. The item's own recorded page range, 87–99, matches every secondary citation already on record exactly, and the scan itself carries the standard JSTOR Early Journal Content header identifying it as freely redistributable for non-commercial purposes. This is a materially stronger standard than appendix A.2's Milnor item, blocked by controlled digital lending on the very same platform, and is closed on the same footing as section 13.7's Borrelli–Jabrane–Lazarus–Thibert paper: the published original itself, read in full, not a translation or a secondary paraphrase.*

**Key Result**

*[What the primary text actually contains, against section 11.5] Reading Hilbert's own text surfaces three points where it diverges from, or adds to, chapter 11's existing secondary-sourced account. First, structure: Hilbert's own proof of the negative-curvature case proceeds through exactly four numbered lemmas about asymptotic curves (no closed asymptotic curve exists; two asymptotic curves through a point meet nowhere else; no asymptotic curve self-intersects; an equal-arc-length offset of an asymptotic curve along its transverse family is itself an asymptotic curve disjoint from the original), not the eight-lemma exponential-map-and-Cartan–Hadamard structure section 11.5 reported from a modern textbook secondary source; Hilbert's own route is more elementary and entirely self-contained within the asymptotic-net construction, with no appeal to Cartan–Hadamard by name. Second, attribution: Hilbert's own paper attributes the quadrilateral area-equals-angle-excess formula this volume calls proposition 11.6 to Dini (1870), Darboux, and Bianchi's textbook treatments, citing them by name in a footnote — Hilbert's own 1901 paper does not mention Hazzidakis at all, so it cannot itself corroborate section 11.4's attribution of the formula to Hazzidakis's 1880 paper, which remains exactly as unconfirmed as section 11.4's own todobox already flagged; the two questions (who proved the formula first, and what Hilbert's paper itself cites) are distinct, and only the second is settled by this pass. Third, the final contradiction: Hilbert computes the same total area two ways — an increasing sequence of asymptotic-coordinate squares, each bounded above by $2\pi$ via the angle-excess formula regardless of size, against the exact hyperbolic-geometry area $2\pi(\cosh r - 1)$ of an ordinary geodesic disc of radius $r$ centred at the same point, which diverges as $r \to \infty$ — rather than the sine-Gordon degeneration argument on $\omega$ that section 11.5's own proof sketch and section 11.5's secondary account both currently give as the mechanism. Both routes are legitimate ways to derive a contradiction from the same underlying facts, but they are not the same argument, and section 11.5's sketch does not currently reflect the primary source's own, comparatively simpler, closing step.*

**Key Result**

*[A second theorem in the same paper, not currently in this volume at all] Hilbert's 1901 paper is in two parts, not one: a second section, headed* Ueber Flächen von positiver constanter Krümmung, *proves that any singularity-free piece of an analytic surface of constant curvature $+1$ has its extremal principal curvature radius attained only on the piece's boundary, never in its interior, unless the piece is (part of) the unit sphere itself — citing H. Liebmann (*Göttinger Nachrichten, *1899, p. 44) for the closed case (no closed analytic surface of constant positive curvature other than the sphere) and supplying an independent proof of the open, bounded-piece case via a maximum-principle argument on the two principal curvature radii, in Hilbert's own hand. This is, in effect, a rigidity theorem for the sphere sitting inside the same paper this chapter cites only for theorem 11.8's negative-curvature half, and it is not reported anywhere in this volume at present. It is recorded here as a specific, precisely located addition for a future pass to fold into chapter 11 — most naturally as a short new section paralleling section 11.7's own $g = 0$ bullet, which currently asserts the sphere's rigidity only in passing and without citing a source for it.*

**Todo — verify primary source**

This pass closes the Hilbert (1901) access item outright but does not rewrite chapter 11 itself to match: section 11.5's proof sketch, section 11.5's eight-lemma account, and section 11.7's unsourced sphere-rigidity remark all still stand exactly as drafted before this pass, now alongside a primary source that would let each of them be corrected or sharpened. section 11.4's Hazzidakis-1880 citation gap is likewise untouched — appendix A.9 narrows what is known about it (Hilbert's own paper does not supply the citation) without closing it. A future pass rewriting chapter 11's proof to follow Hilbert's own four-lemma-plus-two-area-computations route directly, and adding appendix A.9's sphere-rigidity content as new material, is now precisely scoped and does not require any further access work: the full primary text is in hand.

A later pass did carry out that rewrite (sections 11.5 and 11.6), and a further pass after that resolved the sign discrepancy this section's own appendix A.9 first surfaced between Hilbert/Darboux's Area = (angle sum) $- 2\pi$ and proposition 11.6 as the chapter then stated

it: section 11.4 re-derives proposition 11.6 from scratch from this volume's own sine-Gordon equation and corner-angle bookkeeping, verified both symbolically and by high-precision numerical quadrature, and finds the Hilbert/Darboux sign is the one the volume's own conventions force; proposition 11.6 has been restated to match. section 11.4's Hazzidakis-1880 citation gap noted just above remains open, unaffected by that derivation.

## A.10 Three items added by chapter 13

Chapter 13 named three primary sources not previously on this volume's open list. John Nash's 1954 paper, "$C^1$ Isometric Imbeddings" (*Annals of Mathematics*, second series, volume 60, pages 383–396), was flagged as this appendix's single most promising target for a future access pass, on the strength of its canonical status despite falling outside the pre-1923 public-domain window appendices A.8.1 and A.9 have both exploited elsewhere. Nicolaas Kuiper's 1955 two-part paper, "On $C^1$-isometric imbeddings I, II" (*Nederlandse Akademie van Wetenschappen, Proceedings, Series A*, volume 58, equivalently *Indagationes Mathematicae* 17, pages 545–556 and 683–689), was the same citation chapter 12 had already located and left unread.

**Key Result**

*[Both closed — not in the original, but in a specific, identified translation] A later pass closed both items, though not by obtaining the English originals: È. G. Poznyak's 1957 Russian translations of both papers, published together in* Matematika *1, no. 2 (1957), pages 3–16 (Nash) and 17–28 (Kuiper, Part I), and freely hosted by the Steklov Mathematical Institute's Math-Net.Ru portal, have been located, fetched, and read in full. Section 13.5 now reports Nash's own twist-and-stage construction — including section 13.5's simplicial covering of the positive-definite matrix cone, previously guessed at only impressionistically — directly from that reading, and section 13.6 reports Kuiper's one-normal-field stretching-corrugation refinement (section 13.6) the same way, closing the "why does one fewer dimension suffice" question chapter 13's earlier pass had left open (remark 13.8). Section 13.6.1 records what this specifically licenses: a contemporary (1957), editorially-annotated translation by a named mathematician, not the archival English original, with the distinction kept explicit rather than silently upgraded.*

**Todo — verify primary source**

This does not fully close the pair. Kuiper's second 1955 paper (pages 683–689, the non-compact case) has not been obtained in any form, translation or original, and remains open exactly as before. And the standing caveat itself is not nothing: this appendix's own evidentiary standard elsewhere (e.g. appendix A.2's Milnor pagination, cross-confirmed against Milnor's own bibliography) has generally sought the author's own text specifically; a translation, however contemporary and faithful, is one documented step short of that standard, and section 13.6.1 is this appendix's first explicit acknowledgment of a source read in that intermediate form.

Vincent Borrelli, Saïd Jabrane, Francis Lazarus, and Boris Thibert's 2012 paper, "Flat tori in three dimensional space and convex integration" (*Proceedings of the National Academy of Sciences* 109(19), pages 7218–7223), supplies the explicit flat-torus construction section 13.7 originally reported from secondary description only.

**Key Result**

*[Closed, in the published original — not merely narrowed] A further pass located and read this paper in full, in its own published English original rather than a translation: PNAS deposits its Direct Submissions in open full-text repositories, and this pass's copy carries the paper's original pagination, author affiliations, and editorial handling line ("Edited by Yakov Eliashberg, Stanford University, Stanford, CA, and approved March 9, 2012"), confirming it as the published version. Section 13.7 now reports, directly from this reading rather than from any secondary paraphrase, the construction's specific mechanism (section 13.7: a fixed three-direction decomposition of the isometric default, letting the process run as a repeating three-corrugation block rather than Nash's and Kuiper's patch-by-patch bookkeeping), its own numbered Corrugation Theorem (section 13.7: an explicit product-of-rotations formula for the limiting Gauss map, with the paper's own Riesz-product analogy for the resulting fractal structure), and its own reported numerical figures (section 13.7: specific exploding corrugation counts and a* 10.2% *residual length distortion at the truncated, rendered stage, against* 80% *for the naive starting map). This is a materially stronger standard than the translation-mediated closure of the Nash and Kuiper items above: the published original itself, not an intermediary text.*

This closes all three of chapter 13's originally-flagged primary sources at the appendix's own standard of evidence, save one: Kuiper's second 1955 paper (pages 683–689, the non-compact case), which remains entirely unpursued, recorded above.

## A.11 One item added by chapter 12

Chapter 12 names one primary source not previously on this volume's open list, and unlike every item appendices A.9 and A.10 above eventually closed or narrowed to a specific paywall, this one has had its citation itself corrected twice within the chapter's own drafting history before any access attempt began: N. V. Efimov's original theorem was first cited, from Wikipedia's own reference, to *Doklady Akademii Nauk SSSR* 150 (1963), pages 1206–1209; a later pass, reading the full bodies of two arXiv papers that cite Efimov independently of each other, found both instead citing a different, more specific source treated as primary in chapter 12 today: N. V. Efimov, "Generation of singularities on surfaces of negative curvature" (Russian), *Matematicheskii Sbornik*, volume 64 (1964), pages 286–320 — with the 1963 *Doklady* item kept on record as the likely short announcement preceding the full 1964 treatment, a relationship inferred from a common Soviet-era publication pattern rather than confirmed by reading either text. A further pass located, and confirmed directly against the publisher's own table of contents, an English-language translation: *American Mathematical Society Translations*, Series 2, volume 66 (1968), pages 154–190, the seventh of thirteen papers in a miscellaneous AMS-assembled collection. This narrows the question from "does an English translation exist" (now answered: yes, precisely located) to "is it accessible without payment" (still open): the AMS sells the volume as a $265 eBook, no free excerpt of specifically pages 154–190 was located by that pass, and neither the original 1964 Russian text nor this identified English translation has been obtained or read by any pass to date.

**Todo — verify primary source**

This item is narrowed further than chapter 12's first pass left it — a specific, publisher-confirmed English translation now has an exact volume and page range, where before only the Russian original's citation was on record — but is not closed, and unlike appendix A.2's Milnor item (blocked by controlled digital lending) or appendix A.6's White et al. item (blocked by an explicit `needAccess=true` paywall page), no free-mirror search has yet been attempted specifically for this translation volume; a future pass should attempt one before treating the $265 price as a final barrier. Chapter 12 itself already reports the theorem's proof from a secondary source (Mendonça's paper, read in full) at a level this appendix's own standard elsewhere treats as adequate short of the primary text, so closing this item would sharpen citation provenance rather than change anything currently stated about the proof itself.

A further pass took up exactly the free-mirror search the preceding todobox named. It did not find one: no open copy of *American Mathematical Society Translations*, Series 2, volume 66, pages 154–190, was located by this pass, so the $265 price stands exactly as before. The search did, however, surface a citation this volume had not previously recorded, independently repeated across at least six modern papers citing Efimov's theorem (isometric-immersion and negative-curvature literature spanning roughly 2009–2025, each compiled independently of the others): N. V. Efimov, "Impossibility of an isometric imbedding in Euclidean 3-space of certain manifolds with negative Gaussian curvature," *Doklady Akademii Nauk SSSR* 146, pages 296–299 (1962) — a paper distinct from, and one year earlier than, both citations already on record in chapter 12 (the 1963 *Doklady* 150 announcement and the 1964 *Matematicheskii Sbornik* 64 full treatment). Its title, narrower than the 1963 item's ("certain manifolds" rather than the general bounded-away-from-zero statement), together with several of the same modern sources also citing a still-earlier Efimov note (1953, *Doklady* 93, on a one-to-one projection of a negatively-curved surface, and a 1966 *Russian Mathematical Surveys* survey piece), suggests a Efimov's own route to theorem 12.1 may have proceeded through more than the two-paper (1963 announcement, 1964 proof) sequence this volume currently states, possibly a genuine three- or four-paper progression from a restricted 1962 case to the general 1963/1964 result — a historical detail no source consulted by any earlier pass had surfaced, and one this pass narrows but does not resolve: none of the four now-identified Efimov citations (1953, 1962, 1963, 1964) has itself been read, so the exact logical relationship between them (which proves what, and whether 1962 is a genuine special case or an equivalent restatement) remains inferred from titles and citation patterns alone, not confirmed by any primary text.

**Todo — verify primary source**

Two things follow. First, the free-mirror search this pass ran came back empty, so appendix A.11's original todobox stands unresolved on its own terms. Second, and unprompted by anything the original item asked for, this pass adds a new, purely bibliographic sub-question: whether chapter 12's own two-citation history (1963 announcement, 1964 proof) should be revised to a longer sequence beginning at 1953 or 1962. Neither the chapter's mathematical content nor appendix A.11's own access-status assessment changes on the strength of a title and a citation pattern alone; a future pass would need to read at least the 1962 item's own abstract or a detailed secondary description of it, not just its citation, before revising anything this volume currently states.